

Robust Header Compression (ROHC) in Next-Generation Network Processors

David E. Taylor, *Member, IEEE*, Andreas Herkersdorf, *Member, IEEE*, Andreas Döring, and Gero Dittmann

Abstract—Robust Header Compression (ROHC) provides for more efficient use of radio links for wireless communication in a packet switched network. Due to its potential advantages in the wireless access area and the proliferation of network processors in access infrastructure, there exists a need to understand the resource requirements and architectural implications of implementing ROHC in this environment. We present an analysis of the primary functional blocks of ROHC and extract the architectural implications on next-generation network processor design for wireless access. The discussion focuses on memory space and bandwidth dimensioning as well as processing resource budgets. We conclude with an examination of resource consumption and potential performance gains achievable by offloading computationally intensive ROHC functions to application specific hardware assists. We explore the design tradeoffs for hardware assists in the form of reconfigurable hardware, Application-Specific Instruction-set Processors (ASIPs), and Application-Specific Integrated Circuits (ASICs)

Index Terms—ASIC, ASIP, FPGA, hardware assist, header compression, network processor, reconfigurable hardware, ROHC.

I. INTRODUCTION

HEADER compression provides for more efficient use of link bandwidth in a packet switched network by leveraging header field redundancies in packets belonging to the same flow. The key observation that many packet header fields such as source and destination addresses remain constant throughout the duration of a flow while other fields such as sequence numbers change predictably allows header compression techniques to transmit only a few bytes of header information per packet. Typically, reference copies of the full headers are stored at the compression and decompression points in order to reliably communicate and reconstruct original packet headers. Particulars of packet header field classification into static and dynamic sets depend upon the communication protocols and encoding techniques employed by the compression scheme. We provide a brief history of header compression techniques in Section II.

Link efficiency comes at the cost of processing and memory resources in the nodes communicating over the link. While this

may not be a favorable tradeoff in many environments such as communication over optical fiber, it is particularly advantageous for communication over radio links due to their high cost relative to the provided bandwidth [1]. Recently introduced as a new standard by the IETF, Robust Header Compression (ROHC) seeks to provide reliable header compression for efficient use of links with high loss rates [2]. As an example of the potential value of ROHC, consider a mobile handset transmitting voice datagrams using IPv6/UDP/RTP as shown in Fig. 1. In order to achieve a high perceived quality of response, voice datagrams are typically on the order of 20 bytes while IPv6/UDP/RTP headers are on the order of 100 bytes. ROHC achieves typical header sizes of one to four bytes which could reduce the overhead in our example by a factor of 100 and the total bandwidth consumption by a factor of six.

Clearly, ROHC deployment for wireless networking requires implementation at both ends of the radio link. While implementation in handsets, mobile terminals and “appliances” raises many interesting issues, our discussion focuses on the use of ROHC at access and aggregation nodes in wireless networks. As shown in Fig. 2, Base Station Controllers (BSCs) which concentrate multiple links from Base Station Transceivers (BSTs) in cellular networks are primary examples. Industry estimates suggest that BSC link rates may soon reach 2.5 Gb/s, imposing a uni-directional throughput constraint of over 7.8 million packets per second for minimum size packets. For orthogonal ingress and egress processing of our 120 byte voice datagram example, the throughput constraint is over 5.2 million packets per second. In order to understand the required resources of implementing ROHC in this environment, we provide an analysis of the primary functional blocks employed by ROHC in Section IV.

Due to the need for programmability, flexibility, and rapid deployment of new applications, services, and protocols at network edges, access routers and aggregation nodes increasingly employ network processors. Designed for high-speed packet processing, network processor architectures seek to achieve maximum flexibility while meeting the real-time throughput constraints imposed by the supported link rates. As in most systems, flexibility is provided via software programmable processors, either General Purpose Processors (GPPs) or processors with optimized instruction sets. Based on our functional analysis, we extract architectural implications of ROHC implementation on network processor design for the wireless access environment in Section V. A key result is that the processor memory interfaces should be dimensioned to provide an aggregate bandwidth greater than eight times the bandwidth of the supported link in order to support worst case traffic patterns. This poses a challenge for processors designed

Manuscript received October 9, 2002; revised November 4, 2003 and August 23, 2004; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor F. Neri.

D. E. Taylor is with the Applied Research Laboratory, Washington University in Saint Louis, and Exegy Inc., St. Louis, MO 63127 USA (e-mail: dtaylor@exegy.com).

A. Herkersdorf is with the Institute of Integrated Systems, Technical University of Munich, D-80290 Germany (e-mail: a.herkersdorf@ei.tum.de).

A. Döring and G. Dittmann are with the IBM Zurich Research Laboratory, CH-8803 Ruschlikon, Switzerland (e-mail: ado@zurich.ibm.com; ged@zurich.ibm.com).

Digital Object Identifier 10.1109/TNET.2005.852887




Fig. 1. Example scenario of ROHC compression/decompression of IPv6/UDP/RTP headers for communication over a radio link.

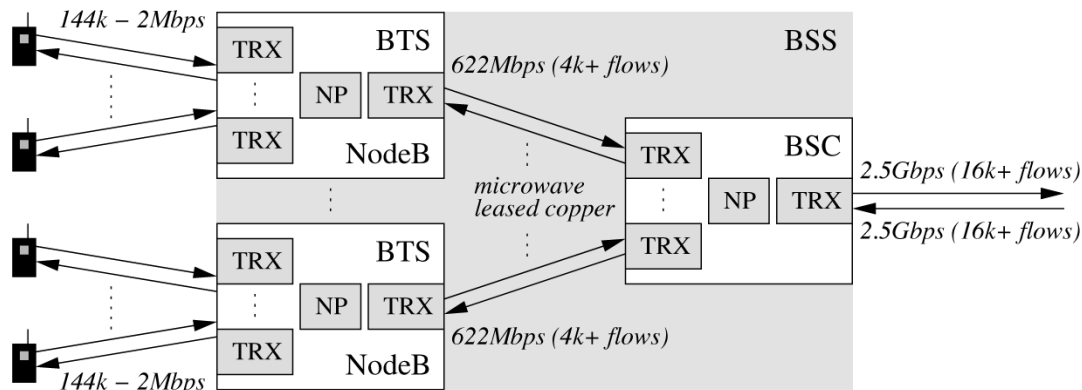


Fig. 2. Block diagram of 3G wireless access network. Link bandwidths are based on future capacity estimates for 3G wireless systems (WCDMA, EDGE, GSM).

to support high speed links or many aggregated links, such as those targeted to the wireless access infrastructure. Based on results from preliminary implementation efforts, ROHC may easily saturate instruction per packet budgets in next generation network processors [3]. In order to meet throughput constraints and maximize the number of available software instructions per packet, many network processors also contain optimized datapaths and application-specific hardware assists for redundant and computationally intensive tasks. While many processor architectures include hardware assists for essential ROHC functions such as packet classification and Cyclic Redundancy Check (CRC) computations, there is no support for ROHC encoding and decoding functions which comprise approximately one third of the per-packet workload. We envision that this will change in the future; hence, we examine the resource requirements and potential performance gains of implementing ROHC encoding and decoding functions in application-specific hardware assists in the form of reconfigurable hardware, Application-Specific Instruction set Processors (ASIPs), and Application-Specific Integrated Circuits (ASICs) in Section VI.

Note that we assume a working knowledge of the IP and UDP protocols. The aforementioned Real-time Transport Protocol (RTP) provides “end-to-end network transport functions suitable for applications transmitting real-time data, such as audio, video or simulation data, over multicast or unicast network services” [4]. For the purpose of our discussion, it is only important to note that each header contains a sequence number (SN) and timestamp (TS) used for maintaining packet ordering and calculating link jitter, respectively.

II. A BRIEF HISTORY OF HEADER COMPRESSION

The notion of header compression was introduced by Jacobson in 1990 for the purpose of improving interactive ter-

restrial response of low-speed serial modem links [5]. Nine years later, Degermark, Nordgren, and Pink developed techniques for compressing permutations of IPv4 headers, IPv6 base and extension headers, TCP and UDP headers, and encapsulated IPv4 and IPv6 headers for transmission over lossy links [6]. That same year Casner and Jacobson added a specification for compressing IP/UDP/RTP headers commonly referred to as CRTP [7].

Seeking to improve performance over lossy links with long round-trip times (RTT), a group of researchers from Ericsson and Lulea University proposed Robust Checksum-based Header Compression (ROCCO) as an Internet Draft in June 1999 [8]. Coupled with transport and link layer protocols capable of partitioning checksum coverage into sensitive and insensitive regions [9], ROCCO was designed to perform well with audio and video codecs which tolerate bit errors in data frames. Studies showed that ROCCO compressed headers to half of the size of those provided by CRTP and remained robust over links with BER rates an order of magnitude higher [10]. In July 2000, Ericsson and Japan Telecom successfully completed a field trial of VoIP over WCDMA using ROCCO [11].

In July 2001, the Robust Header Compression (ROHC) framework was presented as a new standard for header compression by a working group of the IETF [2]. While designed to be a general framework extensible to new protocol stacks, much of the standard focuses on efficiency and robustness of IP/UDP/RTP compression for wireless communication over radio links. Several related standards have been proposed, including a specification for completely removing headers of IP/UDP/RTP packets, termed 0-byte compression, for use over existing cellular links based on GSM and IS-95 [12]–[14]. In May 2002, several companies participated in a successful first trial of the major parts of the ROHC standard including robustness tests over emulated WCDMA/3G links [11].

III. HEADER COMPRESSION WITH ROHC

The general approach of Robust Header Compression (ROHC) is to establish a common context at the compressor and decompressor by transmitting the full packet header, then gradually transition to higher levels of compression via the use of various encoding techniques and packet formats. ROHC is designed to be a flexible header compression framework capable of supporting several protocol stacks. General packet formats, compressor and decompressor finite state machines, modes of operation, error recovery and correction mechanisms, and encoding methods are defined at the framework level. Each supported protocol stack defines a “profile” within the ROHC framework. Profiles fully specify the detailed packet formats, state transition logic, and state actions as well as the encoding methods used for each header field in the protocol stack.

The ROHC framework and profiles for RTP, UDP, ESP, and uncompressed headers are defined in [2]. A specification for running ROHC over the Point-to-Point Protocol (PPP), commonly used for initialization and communication of control information, is defined in [14]. As mentioned in the previous section, a Link Layer Assisted (LLA) profile for IP/UDP/RTP, which is capable of completely removing packet headers, is defined in [13], [12].

For the purpose of our study, we focus on supporting the general ROHC framework as well as the standard IP/UDP/RTP compression profile. RTP/UDP/IP compression operates via the following principle: establish functions from the RTP sequence number (SN) to other fields, then reliably communicate the SN. When functions change, parameters must be updated. The following sections briefly describes the several aspects of the ROHC framework and associated IP/UDP/RTP profile in order to facilitate discussion of our functional analysis and architectural implications on next-generation network processor design. Specifically, we delve into the details of field encoding and decoding techniques in Section III.A as these comprise one third of the processing workload; hence, our analysis of application-specific hardware assists in Section VI focuses on these functions. We refer the reader to our full technical report [16] for a complete overview of ROHC.

A. Field Encoding Methods

The encoding method for dynamic fields in the packet headers is selected based on the dynamic pattern of the specific field. Hence, several encoding methods may be employed for the dynamic fields in a single packet header. Note that while the encoding methods are specified for the ROHC framework, several of the methods are specific to IP/UDP/RTP headers.

1) *Least Significant Bits (LSB) Encoding*: LSB encoding is applied to header fields subject to small changes. Quite simply, the k least significant bits of the field are transmitted. The decompressor derives the original value using a reference value, v_{ref} , which must include the untransmitted bits. Correctness is guaranteed if the compressor and decompressor use interpretation intervals in which the original value resides and the transmitted k LSBs are unique to the original value.

a) *LSB Interpretation Interval Function*: The LSB interpretation interval may be expressed as a function:

$$f(v_{\text{ref}}, k) = [v_{\text{ref}} - p, v_{\text{ref}} + (2^k - 1) - p] \quad (1)$$

For any k , the k LSBs must uniquely identify a value in $f(v_{\text{ref}}, k)$. p is a shifting parameter that may be tuned for efficiency. For example:

- For increasing values, set p to -1
- For proportional shifting of the interval, p may become a function of k ; for example, to have $\frac{3}{4}$ of the interval used for positive changes set p to $2^{k-2} - 1$.

The compressor uses as v_{ref_c} the last compressed value protected by a CRC. The decompressor uses as v_{ref_d} the last decompressed value verified by a CRC. k is selected as the minimum value such that v falls in the interpretation interval $f(v_{\text{ref}_c}, k)$. This selection function is referred to as $g(v_{\text{ref}_c}, v)$.

2) *Window-Based LSB (WLSB) Encoding*: Window-based LSB (WLSB) Encoding provides for robust LSB encoding when the compressor is unable to determine the exact reference value, v_{ref_d} , in use by the decompressor. The compressor maintains a sliding window of possible v_{ref_d} values and calculates k such that all v_{ref_d} candidates produce the correct v . Letting $v_{\text{ref}_{\min}}$ and $v_{\text{ref}_{\max}}$ be the lower and upper bounds of the sliding window, respectively, then k is chosen as follows:

$$k = \max(g(v_{\text{ref}_{\min}}, v), g(v_{\text{ref}_{\max}}, v)) \quad (2)$$

The window is advanced based on positive acknowledgments (ACKs) or the window width limit.

3) *Scaled RTP Timestamp Encoding*: Due to fixed sampling periods employed in real-time applications such as audio and video, the RTP Timestamp (TS) usually increases by an integral of a fixed time interval. When such fixed intervals exist, TS may be derived by the linear function:

$$\text{TS} = \text{TS_SCALED} * \text{TS_STRIDE} + \text{TS_OFFSET} \quad (3)$$

where TS_STRIDE is the fixed time interval, TS_SCALED is the integral multiple, and TS_OFFSET is the linear offset.

TS_STRIDE is explicitly communicated from compressor to decompressor. TS_OFFSET is implicitly communicated by sending several uncompressed TS values from which the decompressor can extract its value. Note that TS_OFFSET is updated locally at TS wraparound via special interpretation interval rules. Following initialization, only TS_SCALED is communicated between compressor and decompressor via WLSB encoding. Note that ROHC also defines a timer-based encoding scheme for the RTP timestamp which we do not address here.

4) *IPv4 Identifier (IP-ID) Offset Encoding*: Offset encoding assumes that the IP-ID increases for each outgoing packet. Therefore, it is more efficient to encode the offset relative to the RTP Sequence Number (RTPSN) where:

$$\text{Offset} = \text{IPID} - \text{RTPSN} \quad (4)$$

For transmission, the offset is compressed and decompressed using WLSB encoding with $p = 0$. Network byte order is synchronized using the NBO or NBO2 flag in the header.

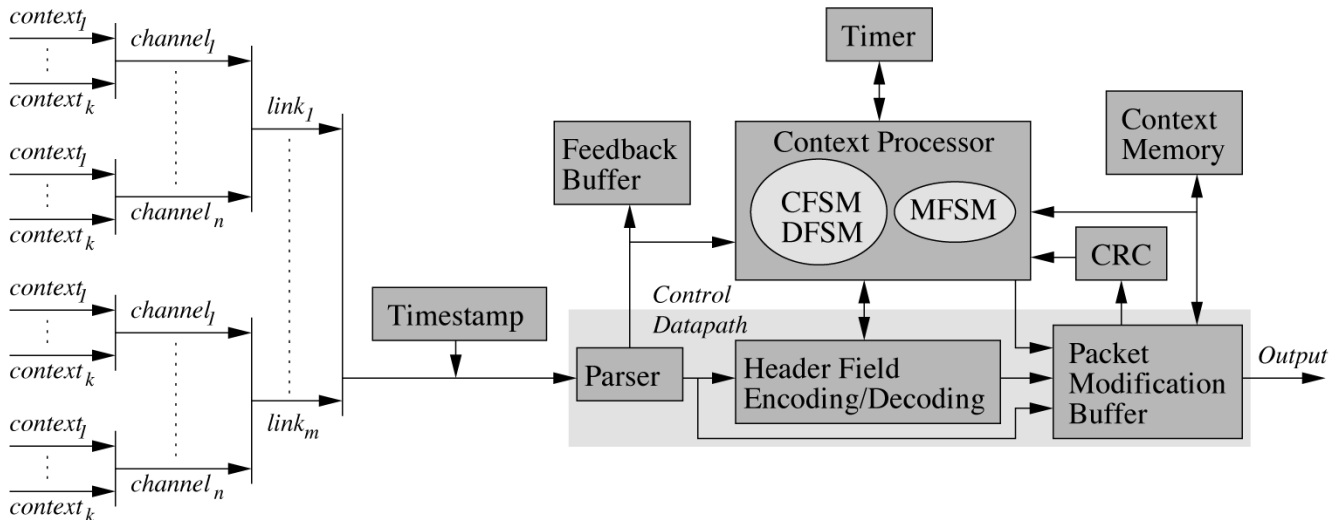


Fig. 3. Logical block diagram of ROHC Compressor/Decompressor for IP/UDP/RTP compression.

B. CRC

For initialization packets, ROHC employs an 8-bit CRC covering all packet fields, excluding the payload. For compressed headers, ROHC calculates the CRC over all of the header fields of the original packet. Header fields are split into two groups, namely CRC-STATIC and CRC-DYNAMIC based on their frequency of change between packets. The static portion of the CRC is calculated over a list of the CRC-STATIC fields concatenated in the order in which they appear in the original packet header, then stored with the context. The CRC calculation continues over the CRC-DYNAMIC fields concatenated in their original order. The static portion of the CRC is only recomputed when CRC-STATIC fields change. Note that the CRC is used to detect bit errors incurred during transmission as well as invalid context information caused by missed or incorrect context updates; therefore, two CRC widths are used to ensure uniqueness between CRC values for context updates.

IV. FUNCTIONAL ANALYSIS OF ROHC

In order to understand the architectural implications of implementing ROHC in a network processor, we must analyze the fundamental operations and interactions of the functional blocks described in the ROHC specification. While there is inevitable variance depending upon implementation decisions and environments, a fundamental analysis provides valuable insight into issues such as resource requirements and scalability. It is important to note that we do not address ROHC performance or compression efficiency, and we attempt to keep our analysis as implementation and traffic independent as possible. Given the absence of open ROHC implementations and packet traces of high bandwidth links with aggregated 3G traffic, we seek to establish useful bounds for network processor architects. To facilitate our discussion without significant loss of generality, we present a logical block diagram of an ROHC Compressor/Decompressor in Fig. 3.

As reflected in Fig. 3, a network processor may support several links. Since there may be multiple channels per link and the number of contexts and compression profile may be configured on a per-channel basis in ROHC, the network processor must

provide a mechanism for identifying the link and channel of all arriving and departing packets. As specified by ROHC, the network processor must also provide the length of link layer frames passed to the compressor/decompressor.

The Packet Modification Buffer is assumed to be a local or dedicated memory space where packets may be temporarily stored during compression and decompression. Context Memory provides storage for all of the per-context information such as state, mode, reference values, etc. The Feedback Buffers allow “piggybacked” feedback to be stored and later retrieved when the context addressed by the feedback is made active via receipt of a packet or specific feedback processing. The remaining blocks are discussed in more detail in the following sub-sections.

A. Context Processor

The Context Processor may be viewed as implementing the core of the ROHC framework. In essence, the Context Processor coordinates packet parsing, context fetches and updates, state and mode transitions, field encoding, and error checking for each ROHC packet. While we assume the general structure and interfaces of the Context Processor remain constant for all ROHC profiles, all packets are processed by three profile-specific sub-blocks: a Packet Parser, a Compression or Decompression Finite State Machine (CFSM/DFSM), and a Mode Finite State Machine. Note that ROHC profiles are assigned on a per-context basis; therefore, the processing performed by each sub-block may change on a per-packet basis for a Context Processor serving many channels.

1) *Packet Parser*: The Packet Parser may be viewed as a finite state machine capable of decoding the many permutations of packet header formats. The minimal set of headers recognized by the parser must include the base and extension headers of the protocols supported by the ROHC profile as well as the profile-specific packet headers. Based on a link and channel identifier, the parser may determine the direction and profile of the packet; for example, ROHC compression for IP/UDP/RTP.

One of the primary functions of the Packet Parser is returning the context identifier (CID) for the packet so that the context in-

formation may be retrieved from Context Memory. For packets with compressed headers, the CID is contained in the ROHC header. The Packet Parser must simply search for the bit-pattern identifying the CID field. For uncompressed headers, the CID must be assigned by a block capable of binding the packet to an established ROHC flow or allocating a new CID for the flow.

The Packet Parser must also extract the header fields requiring encoding or decoding. These fields are passed by the Context Processor to the appropriate encoding or decoding block. When the parser detects “piggybacked” feedback, the information is stored in per-channel, per-context feedback buffers. Finally, the Packet Parser must extract CRC values for error detection.

2) *Compression/Decompression Finite State Machines (CFSM/DFSM)*: Compression and Decompression Finite State Machines (CFSM/DFSM) dictate the format of transmitted packets and feedback. Coupled with the results of packet parsing, field encoding/decoding, and CRC computations, the mode of operation stored with the context information determines state transitions and actions. We performed an analysis to determine the amount of state information required per-context by the CFSM/DFSM [16]. We conclude that the CFSM/DFSM may be implemented as a single profile-specific sub-block of the Context Processor. Each context must simply store the current state and mode, along with the a small set of parameters, totaling 20 bytes of context information.

3) *Mode Finite State Machine (MFSM)*: We assume that a mode variable is maintained for each context. The value of the variable controls context attributes such as state actions and usable packet types. In many cases, the behavior of compressor or decompressor is restricted during mode transitions.

B. Timer

The Timer block shown in Fig. 3 simply denotes the existence of a timing mechanisms available for maintaining timeout timers while in Unidirectional Mode (U-Mode). At minimum, the Context Processor must maintain per-context timeout values relative to a free-running timer. When a packet arrives on a given context, the Context Processor must detect if a timeout has occurred for the context. Several design options exist for such a mechanism and selection of the ideal approach largely depends on the implementation platform.

C. Field Encoding & Decoding

Header field encoding and decoding are major consumers of processing resources. We first present an analysis of the computational requirements of the encoding schemes presented in Section III.A. Memory resources required for encoding parameters are addressed in Section V.F. We performed an analysis of the fundamental operations employed by each encoding and decoding technique in the ROHC specification. We observe that WLSB encoding and LSB decoding are used as building blocks for several encoding techniques. Most encoding techniques contribute only a few additional computations on top of those required for WLSB encoding or LSB decoding. In order to illustrate how a WLSB Encoder may be used as a functional component, a block diagram of an RTPSN Encoder is shown in Fig. 4. Note that shifting parameters are selected based on feedback




Fig. 4. Block diagram of RTP Sequence Number (RTPSN) Encoder employing WLSB Encoder block.

from the WLSB result. This ensures that as the width of the interpretation interval grows beyond a threshold, it is proportionally shifted to provide $\frac{31}{32}$ of the interval for positive changes. Complementary to the encoder, the shifting parameters for the LSB Decoder are selected based on the number of bits received in the ROHC header.

We also observe that the core operations, such as logarithms, employed in LSB encoding are expensive to implement in both software and hardware. Due to the potential usefulness of WLSB Encoder and LSB Decoder blocks, the following sub-sections discuss their design and efficient implementation. We note that the structure of the functions allows for optimizations which eliminate the high cost of implementing the core operations independently.

1) *WLSB Encoder*: As previously stated, WLSB encoding chooses the minimum number of least significant bits of a value to send such that the bits will uniquely identify the original value given a set of likely reference values in use by the decompressor. This task is formalized in (2). The LSB Selection Function employed by the WLSB encoder may be expressed as follows:

$$g(v_{\text{ref}_c}, v) = \min(k | v_{\text{ref}_c} - p \leq v \leq v_{\text{ref}_c} + 2^k - 1 - p) \quad (5)$$

Note that the shifting parameter p may take on values that are functions of k , as specified in the previous examples. In order to find a closed-form solution for k , a parameterized function of p must be set. Based on the previous examples, the following function is proposed.

b) *Parameterizing the LSB Selection Function*: Let p be represented by the function:

$$p = a \times 2^{k-b} - c \quad (6)$$

with passed parameters a , b , and c . The following examples show how p may be set:

- For $p = -1$, set $a = 0, b = 0, c = 1$
- For $p = 2^{k-2} - 1$, set $a = 1, b = 2, c = 1$

Replacing p with the parameterized function, the selection function takes the form of the inequality shown in (7).

$$g(v_{\text{ref}_c}, v) = \min(k | v_{\text{ref}_c} - a2^{k-b} + c \leq v \leq v_{\text{ref}_c} + 2^k - 1 - a2^{k-b} + c) \quad (7)$$

Solving for k yields the inequality shown in (8).

$$g(v_{\text{ref}_c}, v) = \min\left(k \left| \lg\left(\frac{v - v_{\text{ref}_c} - c + 1}{1 - a2^{-b}}\right) \leq k \leq \lg\left(\frac{v_{\text{ref}_c} - v + c}{a}\right) + b \right.\right) \quad (8)$$

Since k must be an integral number of bits, the inequalities and min function may be eliminated as follows:

$$k = \left\lceil \lg \left(\frac{v - v_{\text{ref}_c} - c + 1}{1 - a2^{-b}} \right) \right\rceil \quad (9)$$

Note that the WLSB Encoder performs this computation twice, once with the minimum reference value and once with the maximum reference value, selecting the maximum result. If the value of k computed by the WLSB Encoder is smaller than the minimum value allowed by the available packet format, then the smallest available value of k is used.

c) Efficient Implementation: While (9) may easily be specified in a high-level language for implementation in a General-Purpose Processor (GPP), we continue our analysis to investigate opportunities for optimized implementations. We first note that the negative exponential in the denominator may lead to the need for floating-point division if implemented literally. Logarithms are also typically expensive to implement in software due to the need for approximation arithmetic or lookup tables. In an effort to avoid these expensive computations, we re-formulate the LSB Selection Function as follows:

$$k = \lceil b + \lg(|v - v_{\text{ref}_c} - c + 1|) - \lg(|2^b - a|) \rceil \quad (10)$$

This fnorm of the function eliminates the need for floating point division and reduces the exponential to a simple Logic Shift Left (LSL) operation. The need for an efficient mechanism for computing the binary logarithms remains. Taking advantage of the ceiling operation, the logarithms and remaining arithmetic are efficiently computed using bit string searches.

Let $x = |v - v_{\text{ref}_c} - c + 1|$ and let $y = |2^b - a|$; (10) becomes:

$$k = \lceil b + \lg x - \lg y \rceil \quad (11)$$

We can easily compute the integral part of the binary logarithms by locating the bit-location of the most significant '1'. We refer to the bits to the right of this bit-location as the fractional part. A simple compare of the fractional parts of the logarithms determines whether or not to add an additional bit to the value of k found by adding b and the integral part of $\lg x$ minus the integral part of $\lg y$. To clarify, an example starting from (11) is shown in Fig. 5.

2) LSB Decoder: As previously stated, the decompressor derives the original value from the received LSBs using a reference value, v_{ref_d} , and an interpretation interval. Let m be the number of LSBs received by the compressor. Let $|m|$ be the value of the received LSBs. The original value, v_d , is determined to be the value in the interpretation interval whose m LSBs are equal to $|m|$. Like the WLSB Encoder, the LSB Decoder must employ the parameterized function of p . Hence, the interpretation interval employed by the LSB decoder may be expressed as follows:

$$f(v_{\text{ref}_d}, m) = [v_{\text{ref}_d} - a2^{m-b} + c, v_{\text{ref}_d} + 2^m(1 - a2^{-b}) - 1 + c] \quad (12)$$

Selecting the value within the interval matching the received LSBs may be done in several ways. We focus on an efficient computational implementation which avoids lookup tables or multiple iterations.

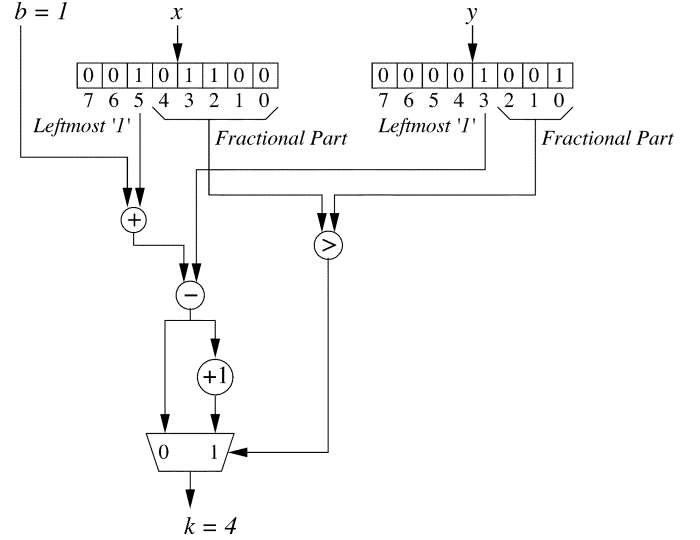


Fig. 5. Example of LSB Selection Function computation via bit string searches. Equation is of the form $k = \lceil b + \lg x - \lg y \rceil$.

d) Efficient Implementation: Let w be the left end-point of the interpretation interval; therefore, $w = v_{\text{ref}_d} - a2^{m-b} + c$. Let $|w|_m$ be the value of the m LSBs of w , which may be expressed as:

$$|w|_m = (w) \text{ modulo } (2^m) \quad (13)$$

Let $|w|_{n-m}$ be the value of the $n-m$ MSBs of w , where n is the width in bits of w . This may be expressed as:

$$|w|_{n-m} = w - |w|_m \quad (14)$$

We note that the original value, v_d , may be selected from the interpretation interval expressed in (12) as follows:

If $|m| \geq |w|_m$, then $v_d = (|w|_{n-m} + |m|)$; else $v_d = (|w|_{n-m} + 2^m + |m|)$.

While this may easily be specified in high-level GPP instructions, we also provide an example of an optimized solution using bit-level optimizations as shown in Fig. 6.

D. CRC

ROHC requires support for CRC computations based on three different polynomials. In the compressor, a CRC checksum is generated from the original header and is transmitted with the header depending on the compression mode. In the decompressor, this checksum is compared to the checksum of the restored header. The computation of CRC checksums with generator polynomials of small degree (here 3, 7, and 8) is comparatively simple. Furthermore, the checksums can be computed on a consecutive sequence of input words. For this type of task look-up based methods are preferred in software implementations and can be used for hardware as well. For a polynomial degree d and input width w bits per lookup, table sizes of $d \times 2^w$ bits result. For software implementations utilizing standard memory widths, it is worth illustrating the performance tradeoffs that arise when deciding how the table entries for polynomial degrees of 3 or 7 bits may be efficiently packed into memory words. For example, utilizing one memory word per 3-bit entry simplifies table address generation but is

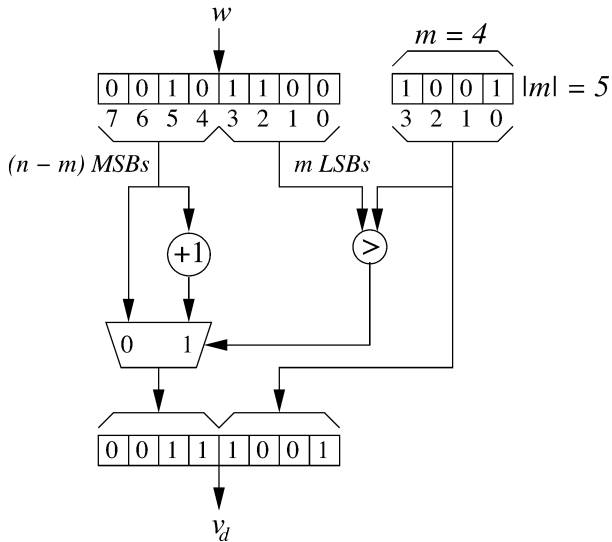


Fig. 6. Example of LSB Decoder operation, where $w = v_{ref,d} - a2^{m-b} + c$ is the left end-point of the interpretation interval and m is the number of received LSBs.

space inefficient. Concatenating 10 3-bit entries and storing them in a single 32-bit word is space efficient but consumes additional processing cycles for table address generation due to the need for an integer divide and shifting operations to extract the desired result from the resulting memory word. Hence, for CRC computations with polynomials of such a small degree, direct computation in software is worth considering. It consists of a mask and parity computation for each output bit. On a processor with a population count instruction, only a few instructions are needed per output bit.

In hardware implementations, the choice of input width of the computation dictates the size and depth of logic. Table I lists the area and logic depth for computing the three polynomials in parallel for input widths of 8, 16, and 32 bits. FPGA area estimates are explained in Section VI.A.1. Note that since the output size is 3, 7, and 8 bits, respectively, the result of the CRC computation fits in one 32-bit word. Hence, all three checksums could be implemented as an instruction set extension and comfortably combined into a single instruction on a 32-bit processor. The individual results could be extracted by shift operations. The same argument holds for a strongly-tailored ASIP implementation.

V. ARCHITECTURAL IMPLICATIONS FOR NETWORK PROCESSORS

Insertion of ROHC compression and decompression in the packet processing path of a network processor presents several architectural implications. Header compression logically resides between the link layer, commonly called Layer 2, and the network layer, commonly called Layer 3, in the protocol stack. Inserting header compression and decompression in the ingress and egress packet processing paths of a network processor requires additional processing, memory, and interconnection resources. It is our intention to establish bounds on the amount of additional resources to serve as “engineering bands” for network processor architects. For the purpose of our analysis, we employ

TABLE I
AREA AND PROPAGATION DELAY FOR CRC CALCULATION. ASIC RESULTS BASED ON AREA OPTIMIZED SYNTHESIS TARGETING IBM Cu-11 PROCESS ASSUMING WORST-CASE DELAYS. FPGA RESULTS BASED ON SYNTHESIS WITH SYNPLICITY SYNPLIFY PRO TARGETING A XILINX VIRTEX-E (–8) DEVICE USING WORST-CASE DELAY ESTIMATES. FPGA AREA ESTIMATES ASSUME 423 LUT/FF PAIRS PER mm^2 .

FSM	Mode	Transition	Parameter(s)	Est. Size
DFSM	All Modes	FC→SC	n_1, k_1, win_1	4 B
	All Modes	SC→NC	n_2, k_2, win_2	4 B
CFSM	U-Mode	IR→FO	$n_{stat}, count_{stat}$	2 B
	U-Mode	FO→SO	$n_{dyn}, count_{dyn}$	2 B
	U-Mode	IR→SO	$n_{sd}, count_{sd}$	2 B
	U/O-Mode	SO→IR	$timeout_1$	2 B
	U/O-Mode	SO→FO	$timeout_2$	2 B
	U/O-Mode	FO→IR	$timeout_3$	2 B
Total				20 B

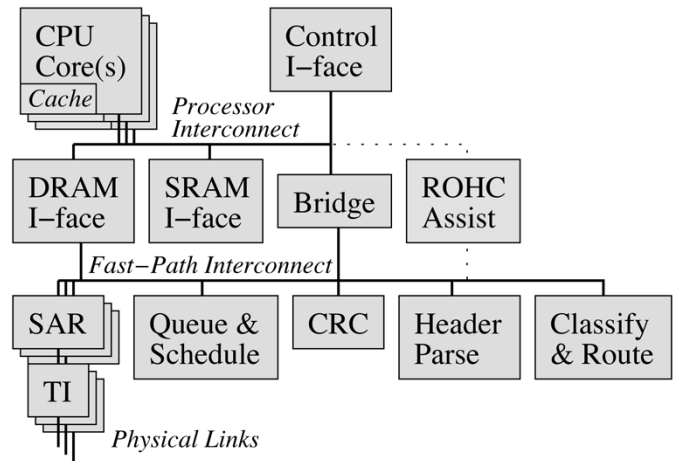


Fig. 7. Block diagram of a generic network processor architecture utilizing hardware assists and hierarchical interconnection.

a generic network processor architecture utilizing hardware assists and hierarchical interconnection as shown in Fig. 7. Note that Fig. 7 includes an ROHC Assist hardware assist block. We motivate the need for this block and analyze the performance and resource tradeoffs in Section VI.

A. Functional Placement

A logic block diagram of the ingress datapath of a network processor supporting ROHC is shown in Fig. 8. Note that following frame reassembly, packets containing compressed headers must be passed to the ROHC Decompressor with an identifier specifying the arrival link and channel as well as the frame length. Header decompression must precede Route Lookup & Classification, as the fields used for packet classification are contained in the original packet header. The network processor must support some form of Internal Packet Route (IPR) function capable of recognizing compressed headers and routing the packets to the ROHC Decompressor prior to classification. This could pose a problem for architectures that implement fixed datapaths between the Frame Reassembly and the Route Lookup & Classification blocks in order to maximize throughput for typical network traffic.




Fig. 8. Logical block diagram of ingress data path including ROHC decompression.




Fig. 9. Logical block diagram of egress data path including ROHC compression.

In order to transmit packets with compressed headers, the network processor must be able to instantiate an ROHC Compressor prior to Frame Segmentation as shown in Fig. 9. As previously mentioned, this requires that an IPR have the ability to identify and pass packets requiring header compression to the ROHC Compressor along with the outgoing link, channel and Context Identifier (CID). Placement of the ROHC Compressor relative to Queuing & Scheduling depends on a number of factors such as the ROHC implementation platform and scheduling algorithms. The following sub-sections discuss the influences and requirements ROHC places on the various components of network processor architecture, beginning with the packet parsing and classification.

B. Frame Reassembly & Packet Parsing

As previously mentioned, high-performance network processors often employ some form of “fast path” in which packets requiring “normal” processing are handled by optimized blocks, bypassing general purpose processing engines. One of the most fundamental and common tasks handled in a “fast path” is link layer frame segmentation and reassembly. ROHC defines a segmentation and reassembly protocol which may require modifications to fixed Frame Segmentation & Reassembly blocks. Note that the protocol is only used when a packet is larger than the largest size supported by the link layer; therefore, its use in many environments is obsolete.

In order to offload the packet classification and route lookup functions, the “fast path” must include some form of header parser which extracts the required header fields. This functionality may be contained in the Frame Segmentation & Reassembly block or implemented as a separate component. In order to leverage “fast path” components for ROHC functionality such as flow classification and CRC computations, the header parser must recognize a new class of header formats defined by each ROHC profile. This suggests that a programmable header parsing block, similar to the one described in [15], would be a favorable architectural feature of a network processor supporting ROHC.

C. Packet Classification

Within the context of ROHC, a flow refers to a sequence of packets that demonstrate the necessary redundancy in static fields as defined by the ROHC profile. In order to bind an uncompressed packet to an ROHC flow, the static fields of the packet headers must be compared against the set of established flows. This is precisely the function performed by

packet classifiers which typically search filter tables used for network management and firewalls. Most packet classifiers are capable of performing filter matches on header fields such as the IP source and destination addresses, transport protocol, and transport source and destination ports. For ROHC compression of IPv6/UDP/RTP and IPv4/UDP/RTP, the fields that must be examined in order to identify a flow are reported in [2]. The total search key sizes for IPv4 and IPv6 headers are 140 bits and 352 bits, respectively. Note that if the IPv6 Flow Label is used (non-zero) the number of bits to be examined could be as low as 20 bits.

At minimum, ROHC requires that the packet classification block support exact match lookups using search keys with configurable widths. This type of search may be efficiently implemented using hashing techniques; however, it is likely that more elaborate algorithms will be employed in packet classifiers employed by next-generation network processors. For high-performance packet classifiers, many network processor platforms make use of Ternary Content Addressable Memory (TCAM) or ASICs implementing proprietary algorithms. We note that the case of binding an IPv6/UDP/RTP header to an ROHC flow could pose a problem for such implementations, as the required 352 bit search key is wider than the maximum width provided by commercially available classifiers.

Ideally, entries in the classification table would contain the CID of the ROHC flow to be compressed. This implies that updates be generated by a block responsible for managing the CID space of each channel. Since there is no explicit flow termination signal in current packet switched networks, a suitable control block must manage the CID space of each channel. This control block could be integrated with the filter and route database manager of the packet classifier. At minimum, the control block must implement a CID allocation algorithm such as Least Recently Used (LRU), manage a set of connection timers, and generate appropriate feedback such as the “REJECT” feedback message used to inform a compressor that no compression context is available for a new flow.

D. Timestamping

A mechanism for assigning an arrival time to ingress packets is required in order to support timer-based RTP Timestamp encoding. Ideally, each packet should be stamped at the transmission interface, immediately following reassembly, and prior to any pre-decompression buffering. The arrival timestamp may be carried in an internal header, or shim, and passed to the decompressor along with the packet. Non-deterministic buffering

delays prior to decompression should be kept to less than a few microseconds. RTP Timestamps are usually on the order of milliseconds; therefore, such small buffer delays should not make contributions to the jitter calculations performed during timer-based RTP Timestamp decoding.

E. Scheduling & Queuing

An in-depth discussion of the affects of ROHC on scheduling and queuing mechanisms in next-generation network processors in large wireless network aggregation nodes is beyond the scope of our discussion. However, we would like to enumerate a few peculiarities in the ROHC standard which imply that further study on this topic would be beneficial.

- 1) Non-deterministic processing time due to radio link properties and application behavior
- 2) Change of total packet length
- 3) Decompressor creates packets in reverse direction for acknowledgments in a bi-directional modes
- 4) Decompressor may create bursts of decompressed packets when using “negative caching”

F. Memory Resources

All header compression schemes achieve transmission efficiency by trading off memory resources. In essence, headers are redundantly stored instead of transmitted. We examine the amount of memory space and bandwidth required to support ROHC in large aggregation nodes.

1) *Space Requirements*: Since the compressor must maintain multiple reference values for sliding windows, there is a significant difference between the amount of memory space utilized by an ROHC compressor and decompressor. In IP/UDP/RTP compression, several fields require the compressor to store multiple reference values in a sliding window when WLSB encoding is used. In order to formulate an upper bound on capacity estimates, we account for per-context memory requirements of a compressor. The memory space required for a reference header and encoding parameters for ROHC compression of IPv4/UDP/RTP and IPv6/UDP/RTP differ by only one byte. Assumptions guiding our choosing sliding windows of width 10 and encoding parameter sizes are provided in the technical report [16]. Using 244 bytes as the upper bound for reference header and encoding parameters and adding the 20 bytes for state machine transition parameters, 264 bytes seems to be a reasonable approximation for the per-context memory requirements for ROHC compression. We also performed this analysis for decompression and found that the per-context memory requirements for ROHC decompression of IPv4/UDP/RTP and IPv6/UDP/RTP are approximately 100 bytes less than that required for compression.

While the specific configuration of contexts per channel, channels per link, and links per port will vary depending on the application or access router configuration, we anticipate that a single network processor may need to support thousands of concurrent contexts. Based on our per-context estimates, total memory requirements for context information exceeds 1 MB for four thousand flows. This implies that off-chip SDRAM should be used for context storage as on-chip memory resources are




Fig. 10. First-order memory bandwidth usage model for ROHC compression in the egress packet processing path of a network processor.

typically limited to a few kilobytes. While commodity SDRAM products provide ample space for context information storage, we examine the additional off-chip memory bandwidth required to support ROHC.

2) *Bandwidth Requirements*: ROHC memory bandwidth consumption will depend heavily upon implementation design decisions and target platforms. Similarly, dimensioning of memory interface bandwidth for network processors is difficult due to the heterogeneity of applications. Processor architects employ various “rules of thumb” in order to gain a first-order approximation of the required memory bandwidth. For the purpose of our analysis, we choose a conservative rule that states the memory interface should provide four times the bandwidth of the aggregate link rate. For example, a network processor supporting an aggregate link throughput of 1 Gb/s should employ a 4 Gb/s memory interface. This rule is based upon the assumption that a packet must be written to memory when received, read from memory for processing, written back to memory after processing, and read from memory for transmission for a total of four transits over the memory interface. It should be noted that some processor architectures employ register pipelines to avoid reading packets from memory for processing. Any packet modifications are applied when the packet is read from memory prior to transmission, resulting in a total of two transits over the memory interface. Given that the wireless access environment requires support for many low-speed or aggregated links, we envision that the number of processing engines in the network processor will be less than the number of links. An architecture employing register pipelines requires an additional packet memory between the ports and pipelines as packets may arrive simultaneously on each link. Since we seek to establish conservative bounds for the additional memory bandwidth required for ROHC, we utilize the single memory interface architecture for our analysis.

In order to gain a first-order approximation for the additional memory bandwidth required for ROHC, we must first establish some assumptions. Due to the large number of supported channels at an aggregation node in the network, we assume that context information is stored off-chip. We also assume that ROHC processing does not require additional reading and writing of packet data, as ROHC processing may utilize an on-chip buffer like the Packet Modification Buffer shown in Fig. 3. Based on these assumptions, the only additional memory accesses generated by ROHC processing is for context fetches and updates.

As shown in Fig. 10, we assume that packet headers are compressed by some factor β that ranges from zero to one. Initialization headers which contain the entire original header correspond to $\beta = 1$, while contexts with high compression efficiency corresponds to $\beta \leq 0.1$. While all of the context information must

be fetched in order to decompress or compress a packet header, only a portion of the context needs to be updated and written back to memory. We make the first-order approximation that the amount of context information written back to memory is proportional to the compression factor β . In general, small headers require few updates of context information while larger headers induce more context information updates.

Based on these assumptions, we find that context accesses for ROHC compression contribute an additional $(1 + \beta) \times C_c$ bytes per packet of memory bandwidth where C_c is the size of the compression context. Similarly, context accesses for ROHC decompression contribute an additional $(1 + \beta) \times C_d$ bytes per packet of memory bandwidth where C_d is the size of the decompression context. As we seek to garner upper-bounds for worst-case design, we will continue our analysis for the compression case since C_c was found to be 264 bytes in the previous section, approximately 100 bytes more than C_d . Note that the analysis for decompression is symmetric with only a change in the value of the context size.

Letting H and P be the length in bytes of the packet header and packet payload, respectively, we assume that packet storage upon arrival and fetch prior to processing consumes $2 \times (P + H)$ bytes per packet of memory bandwidth. Packet storage after processing and header compression and packet fetch for transmission, requires $2 \times (P + \beta H)$. Given a fixed link rate, R , expressed in bytes per second, the number of packets per second equals $\frac{R}{P+H}$. Thus, the amount of memory bandwidth required for a system without ROHC processing is $4 \times R \frac{\text{bytes}}{\text{sec}}$. The expression for memory bandwidth requirements for ROHC compression processing relative to link speed becomes:

$$\text{MemBW} = \left[2 + \frac{264 \times (1 + \beta)}{P + H} + \frac{2 \times (P + \beta H)}{P + H} \right] \times R \frac{\text{bytes}}{\text{sec}} \quad (15)$$

Note that the memory bandwidth scaling factor relative to the link rate is now a function of the header size, payload size, and β . A plot of the memory bandwidth scaling factor versus payload size over the range of β values for ROHC compression of IPv6/UDP/RTP headers is given in Fig. 11. Note that we assumed a static uncompressed header size, H , of 100 bytes.

Supporting our previous example of 20 byte voice datagrams with 100 byte IPv6/UDP/RTP headers would require a total memory bandwidth of approximately 8.4 times the link rate, more than double the bandwidth required for normal packet processing. For a network processor supporting a bi-directional 2.5 Gb/s link or 5 Gb/s total throughput, this implies that the memory interface be dimensioned for 42 Gb/s. Using a standard datapath width of 128 bits, this implies that the interface must be operated at over 328 MHz for standard memory technologies or 164 MHz for Dual Data Rate (DDR) memory technologies [17], [18]. Note that we have assumed simplistic packet handling and application of ROHC compression to 100% of the link traffic. Our results imply that ROHC processing has a significant cost in terms of off-chip memory bandwidth consumption. As link speed increases continue to outpace memory technology improvements, optimized fast-path or header-only processing techniques may become essential to meeting throughput

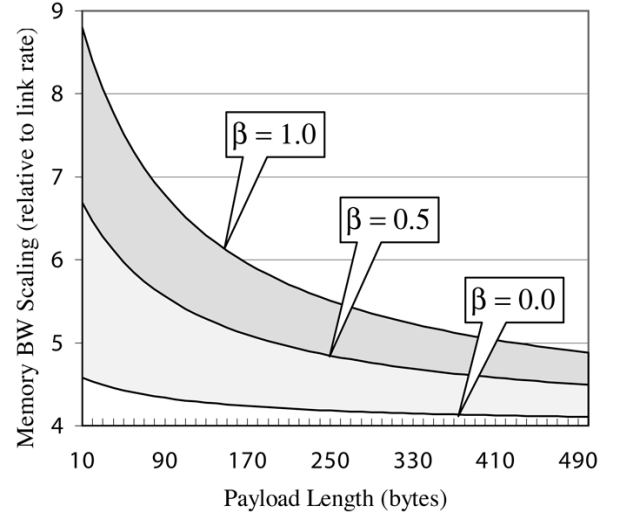


Fig. 11. Memory bandwidth scaling factor relative to supported link rate versus packet payload size for ROHC compression of IPv6/UDP/RTP headers. Assumes static uncompressed header size of 100 bytes.

constraints especially when supporting additional memory-intensive applications.

VI. OFFLOADING ROHC FUNCTIONS TO APPLICATION-SPECIFIC HARDWARE ASSISTS

Based on results from preliminary implementation efforts, ROHC may easily saturate instruction per packet budgets in next generation network processors [3]. In order to meet throughput constraints and maximize the number of available software instructions per packet, many network processors also contain optimized datapaths and application-specific hardware assists for redundant and computationally intensive tasks. In this section, we consider several of the computationally intensive and redundant tasks specific to ROHC as candidates for implementation as hardware assists. A significant portion of the implementation complexity of ROHC lies in the functionality and interactions of the Compressor and Decompressor Finite State Machines (CFSM/DFSM) and the Mode Finite State Machine (MFSM) that establish the context information, encoding parameters, and packet formats. The software-programmable processing engines in network processors are well-suited for such tasks, suggesting a hardware-software co-design to achieve a favorable balance among flexibility and performance. In addition to absolute processing times, the results reported in [3] also provide a breakdown of the amount of time consumed by various steps in the compression process. Correlating these results to the logical block diagram presented in Fig. 3, we make the following observations:

- Packet classification and context binding comprises approximately 14% of the workload
- Packet parsing comprises approximately one third of the workload
- CRC computations comprise 14% to 20% of the workload
- Header field encoding and decoding comprises approximately one third of the workload

As previously discussed, the contributions due to packet parsing, packet classification, and CRC computations may be offloaded to existing hardware assists in current-generation network processor platforms. Offloading the encoding and decoding of header fields allows a majority of the ROHC processing to be done in hardware assists. We begin our analysis of ROHC hardware assists by examining the achievable performance and die area required for Header Field Codecs in several implementation technologies with varying degrees of flexibility. In order to evaluate the implications of a hardware-software co-design, we continue with a discussion of the interconnection bandwidth consumption and hardware-software handovers in Section VI.B. Although we do not provide an explicit analysis, we do believe that reconfigurable hardware assists provide a viable high-performance option for the finite state machines of the Context Processor if additional offloading were required to achieve a specific performance target. As in the previous sections, we focus our analysis on supporting ROHC for IP/UDP/RTP headers.

A. Performance, Flexibility, & Size

In light of our previous analysis of ROHC header field encoding and decoding techniques, we now consider implementation options for a set of ROHC Header Field Codecs hardware assists. Note that we consider two general codec designs: *generic* codecs which require that the shifting parameters be passed with the input values and *field-specific* codecs for each header field which are optimized for the known shifting parameters and field width. Due to the nature of WLSB encoding, all of the encoders may be designed in an *iterative* fashion which seeks to maximize logic reuse and minimize area, or they may be designed in a *pipelined* fashion which computes all intermediate results in parallel and allows a new set of fields to be issued every clock cycle. We implemented the generic and field-specific encoders in both paradigms in order to effectively explore the design space.

We note that given the current lack of insight into ROHC behavior, flexibility is essential for initial implementations. As more experience is garnered, less-flexible implementations providing higher performance may become favorable in the future. We consider three implementation options that represent a likely migration path for ROHC hardware assists. Regarding performance constraints, our analysis assumes a 2.5 Gb/s bi-directional link; therefore, the hardware assists must provide at least 2.6 million header encoding and decoding operations per second. An operation refers to the complete encoding or decoding of all the dynamic header fields of the packet. In order for the generic codecs to meet the throughput constraint, they must provide 7.8 million operations per second (Mops) as they must operate on all three of the dynamic header fields in the IP/UDP/RTP stack. For ease in comparison, performance results are listed with millions of packets per second (Mppts).

1) *Reconfigurable Hardware Assists*: We designed and implemented several codecs in VHDL. The designs were then synthesized targeting a Xilinx Virtex-E device. Achievable clock frequencies are based on worst-case delay estimates in a device with a -8 speed grade. Based on the figures claimed in the Xilinx Virtex-E datasheet [19], a single LUT/FF pair

TABLE II
RECONFIGURABLE HARDWARE RESOURCE USAGE AND PERFORMANCE ESTIMATES FOR ROHC IP/UDP/RTP ENCODING METHODS. LUT/FF USAGE AND CLOCK PERIOD ESTIMATES ACCORDING TO SYNTHESIS WITH SYNPLICITY PRO TARGETING A XILINX VIRTEX-E (-8) DEVICE USING WORST-CASE DELAY ESTIMATES. AREA ESTIMATES ASSUME 423 LUT/FF PAIRS PER mm^2

Width (bits)	ASIC			FPGA		
	Logic Depth	Area mm^2	Tp (ns)	LUTs	Area mm^2	Tp (ns)
8	3	0.0013	1.73	33	.078	2.15
16	4	0.002	1.79	57	.135	3.03
32	4	0.0037	1.87	78	.185	3.42

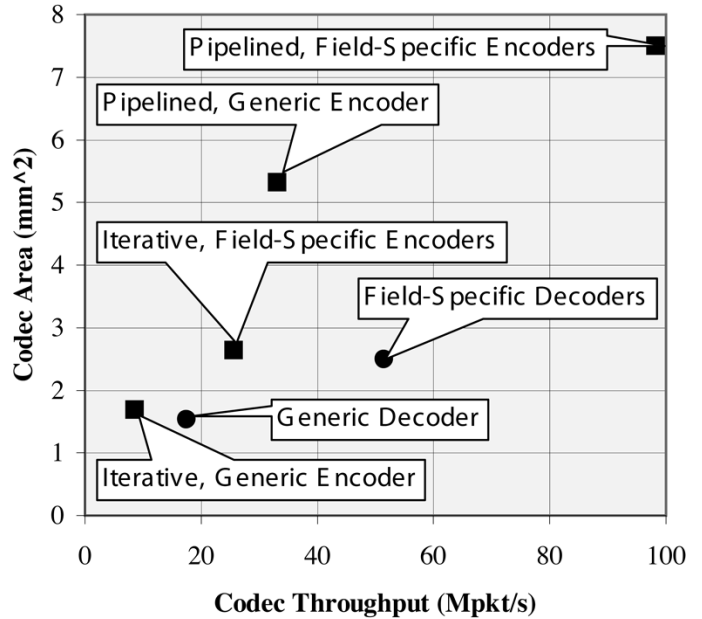


Fig. 12. Area versus throughput tradeoff for various ROHC header field codec designs implemented as Reconfigurable Hardware Assists (RHAs) in embedded FPGA technology.

translates to 13.5 equivalent ASIC gates. A recent announcement by IBM and Xilinx claims that forthcoming embedded Xilinx FPGA cores of 40 k equivalent gates will require 7 mm^2 in the IBM Cu-08 process [20]. We therefore use the estimate of 423 LUT/FF pairs per mm^2 for our area estimate. Our implementation results are shown in Table II.

A fairly large degree of size and speed optimization is achievable in FPGA technology via low-level description and hand-placement of designs; however, our analysis focuses not on determination of absolute performance, but extraction of relative trends between codec designs. Fig. 12 illustrates the throughput and area tradeoffs of employing different types of codecs. In general, *iterative* and *generic* codec designs are more area efficient while a set of *field-specific* and *pipelined* codec designs provided better throughput. For our example case of supporting 2.5 Gb/s links, we note that a combination of an *iterative, generic* WLSB encoder and *generic* LSB Decoder exceeds the necessary throughput of 2.6 million packets per second with better area efficiency than use of field-specific encoders.

For a network processor employing generic codecs or supporting a single ROHC profile, reconfiguration may be

TABLE III
ASIC RESOURCE USAGE AND PERFORMANCE ESTIMATES FOR ROHC IP/UDP/RTP ENCODING METHODS. AREA AND CLOCK PERIOD ESTIMATES ACCORDING TO SYNTHESIS TARGETING IBM Cu-11 PROCESS WITH WORST-CASE DELAY ESTIMATES

Block	cycles/op	LUTs	FFs	Area mm^2	T_{clk} ns	Tput $Mops$	Tput $Mpkts$
Generic LSB Encoder (<i>Pipelined</i>)	1	1077	273	2.546	10.107	98.941	32.908
Generic LSB Encoder (<i>Iterative</i>)	2	694	297	1.641	9.749	51.287	17.096
Generic WLSB Encoder (<i>Pipelined</i>)	1	2250	603	5.319	10.057	99.433	33.144
Generic WLSB Encoder (<i>Iterative</i>)	4	714	417	1.688	9.77	25.589	8.530
IP-ID Encoder (<i>Pipelined</i>)	1	336	187	0.794	7.465	133.958	133.958
IP-ID Encoder (<i>Iterative</i>)	2	193	165	0.456	7.447	67.141	67.141
RTP SN Encoder (<i>Pipelined</i>)	1	559	221	1.322	7.547	132.503	132.503
RTP SN Encoder (<i>Iterative</i>)	4	198	151	0.468	7.379	33.880	33.880
RTP Scaled TS Encoder (<i>Pipelined</i>)	1	2279	590	5.388	10.181	98.222	98.222
RTP Scaled TS Encoder (<i>Iterative</i>)	4	727	386	1.719	9.749	25.644	25.644
Generic LSB Decoder	1	647	210	1.530	19.036	52.532	17.511
IP-ID Decoder	1	188	166	0.444	12.203	81.947	81.947
RTP SN Decoder	1	248	155	0.586	12.082	82.768	82.768
RTP Scaled TS Decoder	1	621	200	1.468	19.407	51.528	51.528

relatively infrequent implying that manually triggered reconfiguration of RHAs would be sufficient. However for network processors supporting multiple ROHC profiles or higher link rates requiring the use of profile-specific codec designs, run-time reconfiguration mechanisms would be required incurring an additional hardware cost. We defer discussion of such mechanisms at this time. Based on our preliminary implementation results, we also predict that implementation of a full ROHC Compressor/Decompressor as an RHA would likely require more die area than available on-chip in the announced embedded FPGA cores. As additional service applications such as content filtering and transcoding are pushed to the network edge, fully offloading ROHC processing may become desirable.

2) *Application-Specific Instruction-Set Processors (ASIPs)*: An ASIP is an alternative between GPPs and ASICs that combines the required flexibility with efficiency by designing a specialized processor core for a class of applications—in our case header compression—such that it is software programmable but its structure is optimized to speed up the common case. Interfaces to the environment of the ASIP, including memories, are tailored to the application class, and processor infrastructure that does not considerably contribute to processing performance is pruned off. This avoids logic that is only rarely used but consumes area and power.

A significant improvement over a GPP stems from a specialized instruction set that speeds up frequently used code sections and reduces program-memory requirements by implementing common combinations of operations in hardware as extended instructions. An example for an ASIP specialized in header parsing for 10 Gb/s is given in [15]. The size of the parser, including a small instruction memory, is in the order of 0.45 mm² in a 0.18 μm-process which demonstrates the area efficiency of the ASIP approach. Examples of methods for the derivation of ASIP instruction sets from application representations can be found in [22], [23].

An ASIP for header compression may employ instructions of different complexity. Single instructions can implement complete coding blocks, such as a WLSB encoder or an LSB decoder, or they can be more basic functions such as a parameterized version of Fig. 5, that can be used to implement a variety of codecs. Herein lies the tradeoff between efficiency, defined

as performance per area, and flexibility. In the ROHC case, the instructions should represent the basic functions that make up a compression profile. When new compression profiles are specified they can be implemented with those basic ROHC functions and added to the ROHC ASIP's functionality. An ASIP designed exclusively for ROHC may implement the static framework around the profiles in a hard-wired fashion. This solution combines the flexibility for new profiles with the most efficient implementation of the framework.

3) *Application-Specific Integrated Circuits (ASICs)*: As the ROHC specification matures and becomes more stable, flexibility may no longer be a necessity and performance may become a higher priority. In such a case, Application-Specific Integrated Circuits (ASICs) may be the preferred implementation technology for ROHC hardware assists as they provide high-performance at low area cost for fixed functions. Results from synthesis of select ROHC codecs targeting the IBM Cu-11 process with worst-case delay estimates are shown in Table III. Note that design tools for ASIC synthesis provide for a wide range of optimization parameters. We provide results for both area and speed optimized synthesis runs in order to illustrate the spectrum of achievable results. Based on these results, a set of generic codecs capable of performing 75 million encodes and decodes per second would require less than 0.1 mm² of die area.

B. Interconnection Architecture

Hardware assists reduce processor cycle consumption at the cost of additional overhead for communication between processor and hardware assists. It is our aim to derive a first-order approximation of the additional interconnection bandwidth required for software/hardware handovers for ROHC hardware assists. This approximation will aid our discussion of interconnection architectures suitable for ROHC hardware assists, including placement and coupling to the processor.

For the purpose of our analysis, we assume a generic processor architecture as shown in Fig. 7 and that the following steps contribute to interconnection bandwidth consumption:

- 1) Packet receive (TI/SAR to SDRAM interface): $P + H$
- 2) Packet load (SDRAM Interface to processor cache): $P+H$
- 3) Context load (SDRAM Interface to processor cache): C_c

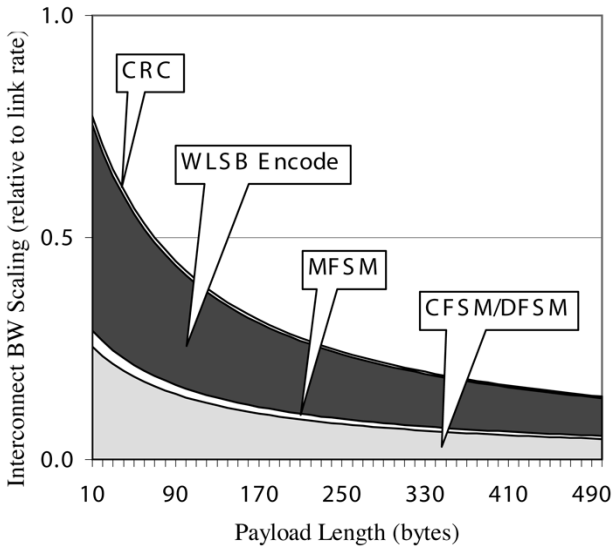


Fig. 13. Additional interconnect bandwidth requirement relative to supported link rate versus packet payload size for ROHC compression of IPv6/UDP/RTP headers with hardware assists. Assumes hardware assists implementing CFSM/DFSM, MFSM, *Generic* WLSB Encoder, and CRC. Assumes static uncompressed header size of 100 bytes and excludes contributions for packet classification, frame CRC, and interconnect overheads such as arbitration and addressing.

- 4) Vector to hardware assists: $HW A_i$
- 5) Return from hardware assists: $HW A_j$
- 6) Context store (Processor cache to SDRAM interface): βC_c
- 7) Packet store (Processor cache to SDRAM interface): $P + \beta H$
- 8) Packet transmit (SDRAM Interface to TI/SAR): $P + \beta H$

The quantities following each step refer to the amount of data per packet which must transit the interconnect. Note that under these assumptions, interconnection bandwidth usage is equal to memory bandwidth usage when no hardware assists are employed. Clearly, these assumptions do not take into consideration the overhead incurred for interconnect transactions. For example, in high-performance bus-based interconnects several cycles are consumed for arbitration and addressing. We assume that such overheads may be accounted for by a general additive constant.

The values of $HW A_i$, the size of arguments passed to the hardware assists, and $HW A_j$, the size of results returned to the processor, depend on the type of ROHC hardware assists employed in the system. We examine the case of ROHC encoding employing a hardware assists for generic WLSB encoding, CRC calculations, MFSM and CFSM state transitions. Based on our previous assumptions, the total amount of information transferred, $HW A_i + HW A_j$, is approximately 85 bytes. Given that the amount of data per packet passed between the processor and hardware assists is a constant, the amount of additional interconnection bandwidth required for these transactions relative to the supported link rate decreases with the packet size. A plot of the additional interconnection bandwidth requirement relative to the supported link rate versus packet payload size for ROHC compression of IP/UDP/RTP headers with hardware assists is shown in Fig. 13 with the contributions for each hardware assist block indicated by a shaded region.

Relative to the interconnection bandwidth consumed by packet and context load and stores, the additional amount required for ROHC hardware assists is small. For our 20 byte voice datagram example an additional $0.7 \times R \frac{\text{bytes}}{\text{sec}}$ of interconnection bandwidth is required. These results imply that ROHC hardware assists may be loosely coupled to the packet processor and utilize standard interfaces to on-chip interconnect. If addressing and arbitration overheads are sufficiently large, hierarchical interconnect may be used to eliminate the contributions of hardware assist communication from the *FastPathInterconnect* as shown in Fig. 7.

VII. CONCLUSION

We have provided an overview of Robust Header Compression (ROHC) and extracted the primary functional blocks required for an ROHC Compressor/Decompressor. From this analysis we examined the architectural implications imposed by ROHC on network processors designed for use in wireless access infrastructure such as Base Station Subsystems (BSS). Based on reasonable assumptions regarding the probable environment of use we provided an estimate for the amount of memory required to store context information. Assuming a context size of 264 bytes per context and a network processor supporting thousands of flows implies that off-chip memory must be used for context storage. We then analyzed the required memory bandwidth relative to the supported link rate and found that for small payload sizes ROHC requires a memory interface providing a bandwidth of up to nine times the link rate. This is a significant result given standard industry practice of dimensioning network processor aggregate memory bandwidths equal to the link rate or at most four times the link rate.

Based on initial implementation results and capacities of existing network processors, we argue that ROHC imposes a significant strain on processing resources. We then explored the design space for application-specific hardware assists for ROHC in the form of reconfigurable hardware, Application-Specific Instruction-set Processors (ASIPs), and Application-Specific Integrated Circuits (ASICs). We showed that the additional interconnection bandwidth required for software/hardware handovers is relatively small implying that hardware assists could be loosely coupled to the packet processor and employ standard interfaces to on-chip interconnect. We assert that supporting ROHC in next-generation network processors targeted to large aggregation nodes in wireless access networks requires consideration at the architectural level. We also provide evidence for continued incorporation of application-specific hardware assists in high-performance network processor architectures.

ACKNOWLEDGMENT

The authors would like to thank M. West for his prompt replies to their questions regarding his slides presented at the ROHC Working Group meeting of the IETF. They also would like to thank P. Buchmann for his assistance with CAD tool flows for the ASIC hardware assists evaluation.

REFERENCES

- [1] "Effnet: An Introduction to EffnetEdge Header Compression Technology," Effnet Inc., Mountain View, CA, 2002.
- [2] C. Bormann *et al.*, "RObust Header Compression (ROHC): Framework and Four Profiles: RTP, UDP, ESP, and Uncompressed," IETF Network Working Group, RFC 3095, Jul. 2001.
- [3] M. West, "ROHC implementation experience," presented at the ROHC Working Group Meeting, 50th IETF Meeting, Minneapolis, MN, Mar. 2001.
- [4] H. Schulzrinne *et al.*, "RTP: A Transport Protocol for Real-Time Applications," IETF Network Working Group, RFC 1889, Jan. 1996.
- [5] V. Jacobson, "Compressing TCP/IP Headers for Low-Speed Serial Links," IETF Network Working Group, RFC 1144, Feb. 1990.
- [6] M. Degermark, B. Nordgren, and S. Pink, "IP Header Compression," IETF Network Working Group, RFC 2507, Feb. 1999.
- [7] S. Casner and V. Jacobson, "Compressing IP/UDP/RTP Headers for Low-Speed Serial Links," IETF Network Working Group, RFC 2508, Feb. 1999.
- [8] L.-E. Jonsson, M. Degermark, H. Hannu, and K. Svanbro, "RObust checksum-based header compression (ROCCO)," IETF Network Working Group, Internet Draft, Jun. 1999.
- [9] L.-A. Larzon, M. Degermark, and S. Pink, "Efficient Use of Wireless Bandwidth for Multimedia Applications," Lulea Univ. of Technology, Lulea, Sweden, Tech. Rep., 2000.
- [10] L.-A. Larzon, H. Hannu, L.-E. Jonsson, and K. Svanbro, "Efficient transport of Voice over IP over cellular links," in *Proc. IEEE Globecom*, 2000, pp. 1669–1676.
- [11] M. Thoren, "Ericsson Compression Technology Boosts 3G Network Capacity," Ericsson Press Release, May 2002.
- [12] L.-E. Jonsson, "RObust Header Compression (ROHC): Requirements and Assumptions for 0-byte IP/UDP/RTP Compression," IETF Network Working Group, RFC 3243, Apr. 2002.
- [13] L.-E. Jonsson and G. Pelletier, "RObust Header Compression (ROHC): A Link-Layer Assisted Profile for IP/UDP/RTP," IETF Network Working Group, RFC 3242, Apr. 2002.
- [14] C. Bormann, "RObust Header Compression (ROHC) over PPP," IETF Network Working Group, RFC 3241, Apr. 2002.
- [15] G. Dittmann. (2000) Programmable Finite State Machines for High-Speed Communication Components, Master's Thesis. Darmstadt Univ. of Technology. [Online] http://www.zurich.ibm.com/ged/Header-Parser_Dittmann.pdf
- [16] D. E. Taylor, A. Herkersdorf, A. Döring, and G. Dittmann, "Robust Header Compression (ROHC) in Next-Generation Network Processors," Dept. Computer Science and Engineering, Washington Univ., St. Louis, MO, Tech. Rep. WUCSE-2004-99, Sep. 2002.
- [17] "36Mb DDR SIO SRAM 2-Word Burst," Micron Technology Inc., Datasheet, Dec. 2002.
- [18] "256 Mb Double Data Rate (DDR) SDRAM," Micron Technology Inc., Datasheet, Oct. 2002.
- [19] "Virtex-E 1.8 V Field Programmable Gate Arrays," Xilinx Inc., Xilinx Datasheet (DS022), 2001.
- [20] "IBM, Xilinx shake up art of chip design with new custom product, IBM and Xilinx, Press Release," IBM, <http://www-3.ibm.com/chips/products/asics/products/cores/efpga.html>, Jun. 2002.
- [21] D. E. Taylor, J. S. Turner, J. W. Lockwood, and E. L. Horta, "Dynamic hardware plugins (DHP): Exploiting reconfigurable hardware for high-performance programmable routers," *Computer Networks*, vol. 38, pp. 295–310, Feb. 2002.
- [22] M. Arnold and H. Corporaal, "Designing domain-specific processors," in *Proc. 9th Int. Symp. Hardware/Software Codesign (CODES'01)*, Apr. 2001, pp. 61–66.
- [23] I.-J. Huang and A. M. Despain, "Generating instruction sets and microarchitectures from applications," in *Proc. Int. Conf. Computer Aided Design (ICCAD'94)*, Nov. 1994, pp. 391–396.




David E. Taylor (M'04) is currently developing high-performance reconfigurable hardware systems at Exegy Inc. Prior to joining Exegy, he was a Visiting Assistant Professor in the Department of Computer Science and Engineering and was actively involved in computer communications research at the Applied Research Laboratory at Washington University in Saint Louis. He received the Doctor of Science degree in computer engineering in August 2004, M.S. degrees in electrical and computer engineering in May 2002, and B.S. degrees in electrical

and computer engineering in December 1998 from Washington University in Saint Louis. His research interests include the design and analysis of scalable searching algorithms and architectures, IP lookup and packet classification algorithms, high-performance reconfigurable hardware systems, programmable routers, and network processors. He held a Research Internship with the network processor hardware group at the IBM Zurich Research Laboratory during the summer of 2002. Dr. Taylor has been a member of the ACM since 2004.



Andreas Herkersdorf (M'91) is a Full Professor in the Department of Electrical Engineering and Information Technology at the Technical University of Munich (TUM), Germany. He received a Dipl.-Ing. degree in electrical engineering from TUM in 1987, and a Ph.D., also in electrical engineering, from the Swiss Federal Institute of Technology (ETH), Zurich, in 1991. Between 1988 and 2003 he was with the IBM Zurich Research Laboratory, Ruschlikon, Switzerland. In September 2003 he became head of the Institute of Integrated Systems at TUM. His

research interests include reconfigurable multi-processor VLSI architectures for IP networking and automotive applications, system level SoC modeling and design space exploration methods, and autonomic computing.



Andreas C. Döring is a Research Staff Member at the IBM Zurich Research Laboratory. He received his Ph.D. from the University of Luebeck (2001), and his Diploma degree in computer science from the University of Karlsruhe in 1995. His current research topic is address translation for server I/O. Further research interests include filesystems, application-specific reconfigurable circuits, fault-tolerant, and parallel computing.



Gero Dittmann received a Dipl.-Ing. (M.Sc.) degree in electrical engineering from Darmstadt University of Technology, Germany, in 2000. In the same year, he joined the IBM Zurich Research Laboratory in Switzerland where he is working as a researcher in the I/O Network Architecture group. Furthermore, he is pursuing a Ph.D. at the Computer Engineering and Networks Laboratory (TIK) of the Swiss Federal Institute of Technology (ETH) in Zurich. His current research interests include ASIP design methods, communication networks, and interconnect architectures for high-performance computing systems.

tures for high-performance computing systems.