

Real time Protocols

Michael Kropfberger, 9555885
mailto:michael.kropfberger@gmx.net

Seminar of Mathematics and Theoretical Computer Science

Prof. Dr. Patrick Horster

(620.710 WS 00/01)

Sensitization

- demand for multi-media services
 - live reports
 - video on demand (VoD)
 - voice over IP (VoIP)
 - audio/video/whiteboard conferences
- HTTP and FTP → not prepared (no QoS, no streaming)

Existing IETF Protocols

- Real-Time Protocol (RTP)
- Real-Time Control Protocol (RTCP)
- RTP Profiles
 - Audio and Video Conferences
 - Secure RTP
- Resource ReSerVation Protocol (RSVP)
- Real-Time Streaming Protocol (RTSP)

Real-Time Protocol

- runs on ST-II, UDP/IP, IPX or ATM AAL5
- UDP/IP widely used
 - unreliable
 - multicasting
- RTP inherits UDP “features” \rightsquigarrow RTCP

Mixers vs. Translators

- Mixers
 - change the stream, eg. from high to low quality
 - generates new source identifier (SSRC)
 - keeps track of originating, contributing sources (CSRC)
- Translators
 - eg. firewalls: translate from multicast to unicast
 - translate from UDP/IP to ATM AAL5
 - forward RTP packets, keep SSRC intact

RTP/UDP/IP Header

Version	Hdr Lngth	ToS	Length(bytes)			
Identification			Flag	FrgmOffs		
Time to Live		Protocol	Header checksum			
Source IP Address						
Destination IP Address						
Options (if any)						
Source Port				Dest Port		
Datagram Length				Checksum		
Vers	Pad	eXt	CC	Mark	Pay	SequenceNr
Timestamp						
Synchronization Source Identifier						
(first) Contributing Source Identifier						
(other) Contributing Source Identifier						
(last) Contributing Source Identifier						
Profile-Specific Information						

IPv4-Header

UDP

RTP

RTP/UDP/IP Header: Details

- Payload Type (7 bits)
 - according to the profile
 - eg. GSM, MPEG-1 layer 3, JPEG, MPEG-2 video, H.261...
- Sequence Number (16 bits)
 - detect packet loss
 - restore packet sequence
 - randomly initialized (if not for server, then for possibly encrypting translators)

RTP/UDP/IP Header: Details 2

- Timestamp (32 bits)
 - monotonically and linearly clock
 - might be differently ordered like MPEG-2 I,B and P frames
 - clock frequency \gg sample rate
- Synchronization Source Identifier (SSRC) (32 bit)
- Contributing Source Identifiers (CSRC) (32 bit)

RTP/UDP/IP Header Compression

- 40 byte per RTP/UDP/IP header
- 20 ms packetization interval \rightsquigarrow 16 kbit/s only for headers
- serial line compression down to 2 bytes (4 bytes with checksum)
- how?
 - only 50% of the fields change
 - other changes in a predictable way
 - generate 8 bit CIDs, store first packet and first-order difference
 - fall-back to full packets on other changes

Real-Time Control Protocol

- four tasks
 - send information about QoS (lost packets, jitter...)
 - transfer clear-text information (eg. CNAME)
 - calculate RTCP packetization rates
 - keep track of all joined participants
- packet types
 - sender
 - receiver
 - SDES
 - BYE

V		Report Cnt	Ptype:200	Length
SSRC of Sender				
NTP Timestamp				
RTP Timestamp				
Sender's Packet Count				
Sender's Byte Count				
SSRC of first source				
% Lost		Cummulative Packets Lost		
Extended Highest Sequence Number Received				
Interarrival Jitter				
Time of last Sender Report				
Time since Last Sender Report				
..List of Sender Reports				
SSRC of last source				
% Lost		Cummulative Packets Lost		
Extended Highest Sequence Number Received				
Interarrival Jitter				
Time of last Sender Report				
Time since Last Sender Report				
Profile-specific Information				

V		R Cnt	Ptype:201	Length
SSRC of Sender				
SSRC of first source				
% Lost		Cummulative Packets Lost		
Extended Highest Sequence Number Received				
Interarrival Jitter				
Time of last Sender Report				
Time since Last Sender Report				
..List of Sender Reports				
SSRC of last source				
% Lost		Cummulative Packets Lost		
Extended Highest Sequence Number Received				
Interarrival Jitter				
Time of last Sender Report				
Time since Last Sender Report				
Profile-specific Information				

Y	R Cnt	Ptype:202	Length
SSRC/CSRC of first source			
SDES items			
further SDES items			
.....			
... List of other SSRC/SDES chunks			
SSRC/CSRC of last source			
SDES items			
further SDES items			
.....			

Where are we?

- Real-Time Protocol (RTP)
- Real-Time Control Protocol (RTCP)
- RTP Profiles
 - Audio and Video Conferences
 - Secure RTP
- Resource ReSerVation Protocol (RSVP)
- Real-Time Streaming Protocol (RTSP)

RTP Profile: Audio and Video Conferences

- 20 ms packetization rate (or the formats natural frame size)
- sampling frequency out of 8000, 11025, 16000, 22050, 24000, 32000, 44100, 48000
- channel ordering from left-to-right
- audio encodings
 - eg. G.722, G.723, G.726, G.728, G.729, GSM, MPA, RED
- video encodings
 - Motion JPEG, H.261, H.263, MPV


Where are we?

- Real-Time Protocol (RTP)
- Real-Time Control Protocol (RTCP)
- RTP Profiles
 - Audio and Video Conferences
 - Secure RTP
- Resource ReSerVation Protocol (RSVP)
- Real-Time Streaming Protocol (RTSP)

RTP Profile: Secure RTP

- provides privacy, message authentication, and replay protection
- additive AES compliant stream cypher in counter mode
- MAC over the whole packet
- contains data according to another profile
- normal RTP header format, adds a 4-byte authentication tag

Secure RTP Header



Cryptographic Context

- encryption key k_e (fixed for session)
- message authentication key k_m (fixed for session)
- 32-bit rollover counter r (which counts how many times the 16-bit *RTP sequence number* wrapped around 0xFFFF)
- the last authenticated sequence number s_l
- replay list L (only receiver side), keeps track of already processed packets

Cryptographic Background

- AES compliant symmetric key block-cypher
 - Advanced Encryption Standard
 - 128 bit block size
 - three key sizes of 128, 192, 256 bits
 - since October 2000: Winner is “Rijndael”

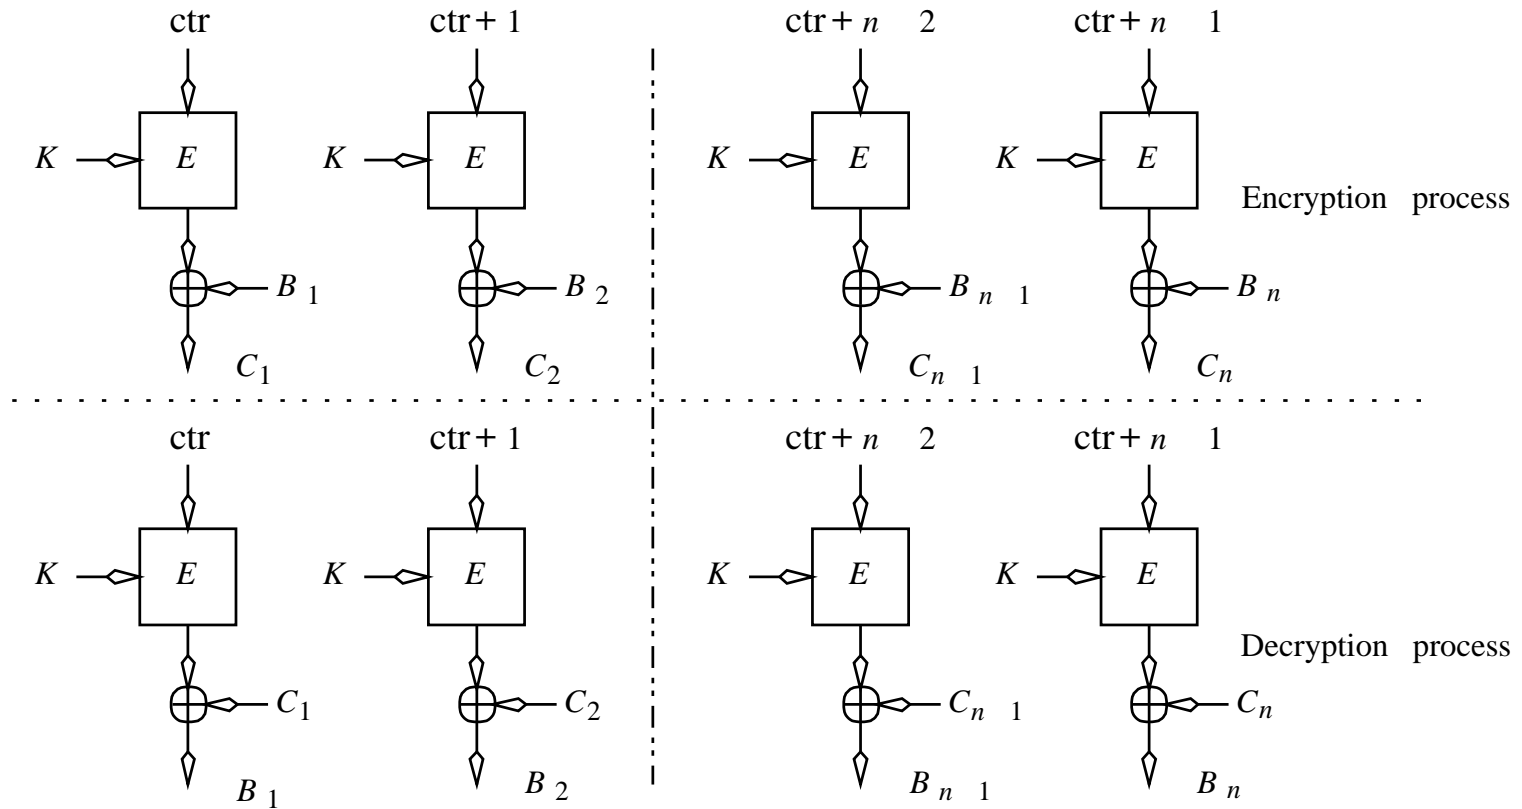
Cryptographic Background: Rijndael

- data block is partitioned into an array of bytes
- each cypher operation is byte oriented
- multiple rounds (10, 12 or 14; depends on key size)
- one round consists of four layers
 - first layer: 8x8 S-Box applied to each byte
 - second and third: shifting of array rows, mixing columns
 - fourth layer: subkey bytes XORed with each byte of array

Cryptographic Background: Counter Mode

- resistancy to redundant plaintext attacks
- cypher block chaining (CBC) mode
 - encrypt last block with key and XOR with new block
 - dependancy on all packets
- segmented integer counter mode (SIC)
 - encrypt counter with key and XOR with new block
 - small hamming distance for ctr and $ctr + 1 \rightsquigarrow$ only problematic for differentially weak ciphers

Cryptographic Background: SIC



Cryptographic Background: SIC in SRTP

- $ctr = [(r * 65536) + seq] * 4096 + i$
- seq is the RTP sequence number, i is a counter for each 128 bit block
- IP packet size = 64 KB \rightsquigarrow 4096 blocks
- Jumboframes are unlikely to be used for multimedia traffic
- ctr has to be unique over the session life time
- maximum of $2^{48} = 281,474,976,710,656$ SRTP packets
- 20 ms packetization time \rightsquigarrow 178,510 years

Cryptographic Background: UMAC

- Message Authentication Code (MAC) → UMAC
- fast (eg. one clock cycle for one byte)
- extra security by key and a “nonce” (our counter ctr)
- $UMAC = E_{AES}(k_m, ctr) \oplus UHASH(k_m, B_i)$
- UMAC-OUTPUT-LEN = 4 bytes

Cryptographic Background: UHASH

- three layers
 - bulk hash function $NH \rightsquigarrow$ speed optimized, compresses block
 - polynomial hash \rightsquigarrow 16 byte
 - inner-product hash \rightsquigarrow 2 byte
- repeat layers with slightly different keys to get additional 2 bytes
- repetition is independent \rightsquigarrow trade authenticity with speed
- allows quick processing to survive DoS attacks

Replay Detection

- bit field over the last `SRTP_WINDOW_SIZE`
- $SRTPseq - SRTP_WINDOW_SIZE$ and $SRTPseq$
- $SRTPseq = r * 65,536 + seq$
- log all older packets and replayed packets (paranoia?)

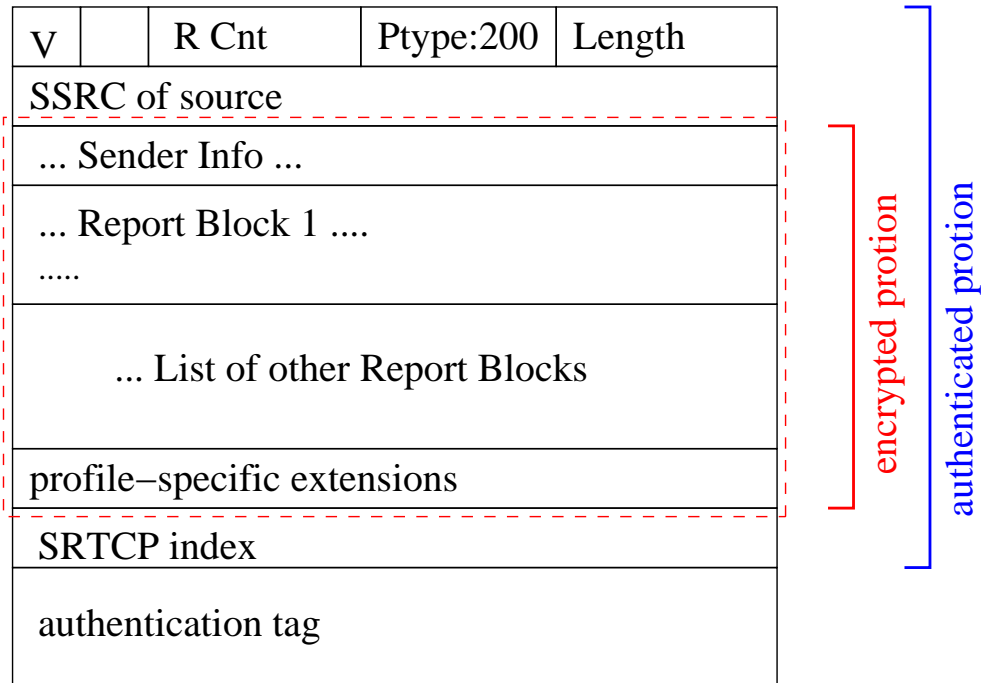
Replay Detection: SRTP WINDOW SIZE

- how big should SRTP_WINDOW_SIZE be?
 - receiver buffer size → “real” real time
 - unidirectional com (eg. radio) → 3 sec
 - video conferences → not more than 400 ms
 - 20 ms packetization rate \rightsquigarrow 50 packets/sec
 - 500 ms buffer \rightsquigarrow max. 25 packets (opt. 32 packets)

Secure RTCP

- based on same ideas as for SRTP
- 32 bit authentication tag
- 32 bit SRTCP index
 - $2^{32} = 4,294,967,296$ packets for one session
 - has to be terminated with RTCP BYE packet

Secure RTCP Header



Where are we?

- Real-Time Protocol (RTP)
- Real-Time Control Protocol (RTCP)
- RTP Profiles
 - Audio and Video Conferences
 - Secure RTP
- Resource ReSerVation Protocol (RSVP)
- Real-Time Streaming Protocol (RTSP)

Real-Time Resource ReSerVation Protocol (RSVP)

- IP: best effort
- RSVP adds rate-sensitive and delay-sensitive QoS
- soft state over routers
- tunneling for non-RSVP networks

Real-Time Streaming Protocol (RTSP)

- may use RTP
- “VCR-style” remote control functionality
- similar syntax to HTTP/1.1
- typical session: DESCRIBE → SETUP → PLAY → PAUSE → TEAR-DOWN

Conclusion

- what did we discuss?
 - Real-Time Protocol (RTP)
 - Real-Time Control Protocol (RTCP)
 - RTP Profiles
 - * Audio and Video Conferences
 - * Secure RTP
 - Resource ReSerVation Protocol (RSVP)
 - Real-Time Streaming Protocol (RTSP)
- ∃ base for real-time streaming of multi-media data, not widely used
 - ⊕ RealNetworks uses RTP, but favors RDP
 - ⊕ RSVP is supported by modern routers (eg. Cisco)

