IEEE P802.15
Wireless Personal Area Networks

Project	IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs)
Title	**Rolling Shutter – Frequency Shift Keying, Input for 802.15.7r1 Draft D0**
Date Submitted	[11 March, 2016]
Source	[Hsin-Mu (Michael) Tsai (NTU), Yen-Ting Liu (NTU), Yu-Lin Wei (NTU), Kate Ching-Ju Lin (Academia Sinica)] [No. 1, Sec. 4, Roosevelt Rd., Taipei 10617, Taiwan] Voice: [+886 2 3366-4888] Fax: [ ] E-mail: [hsinmu@csie.ntu.edu.tw]
Re:	
Abstract	[This is the rolling shutter – frequency shift keying input to the 802.15.7r1 draft, prepared by the team at National Taiwan University.]
Purpose	[To provide text input for draft D0.]
Notice	
Release	

1.0 PHY Layer Operating mode(s)

Optical Camera Communications is introducing three new operating modes.

- PHY 4 accommodates Rolling/Global Shutter Cameras and Low Rate PD
- PHY 5 accommodates Rolling Shutter Cameras
- PHY 6 accommodates 2 Dimensional Screen Codes

See table 73, 74 or 75 from IEEE802.15.7-2011 for an example table.

PHY Operating Modes				
Modulation	**RLL Code**	**Optical Clock Rate**	**FEC**	**Data Rate**
RS-FSK	None	Variable – determined by the set of supported camera receivers	XOR FEC with variable density determined by the supported camera receivers	Variable and determined by the user

2.0 PHY specifications

See IEEE802.15.7-2011 sections 10, 11 or 12 for current specifications.

2.1 Rolling Shutter – Frequency Shift (RS-FSK) Transmitted Waveform

RS-FSK takes advantage of the rolling shutter sampling mechanism in the optical camera receiver, and therefore the "rolling shutter" prefix in the name. However, from the perspective of the transmitter, RS-FSK uses simple frequency shift keying (FSK) signal format. Firstly, a number of K frequencies are used to represent a bit pattern of $\log_2 K$ bits. Secondly, the transmitter uses "square wave pulse shaping", i.e., it will only use two levels, an ON level and an OFF level. This allows us to avoids complex driving circuitry, in particular, a digital to analog converter (DAC), and reduces the cost of the transmitter. Moreover, it will also allow a clean stripe pattern observed in the captured image at the receiver side, which is utilized by the demodulation process. (Please see section 2.3 for a description of the relationship between the width of the strips observed in the pattern in the camera captured image)

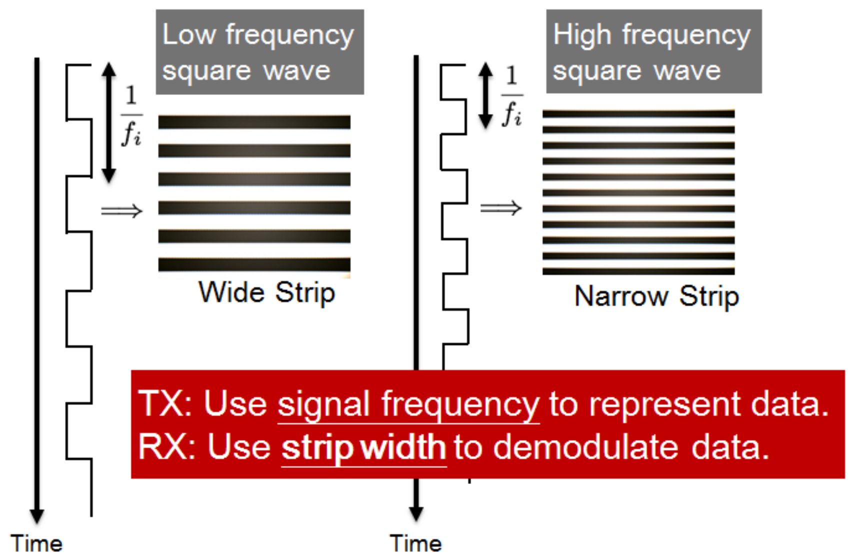

Figure AA – An example of the transmitted RS-FSK signal waveform and corresponding camera captured image

2.1 Transmitted signal frequency

With the exception of the frequency used by the preamble, which is used to detect the start of a PPDU by the receiver, the actual set of frequencies used for data transmission is left to be determined by the user, based on the read-out time of the supported optical camera receiver (see section 2.3). It can be specified by the optical field of the PPDU (see section 4).

RS-FSK specifies the preamble frequency at 2.2 KHz. The majority of the commercially available image sensors have a read-out time of 20 to 30 microseconds. With this preamble frequency, the width of a pair of bright and dark strips in the observed stripe pattern in the camera captured image is between 15 and 23 pixels (see section 2.3 for the method to calculate this), and can be reliably detected. In addition, the obtained from the reception of the preamble frequency can be used to calibrate the value of the read-out duration, allowing better reception error performance.

For the set of frequencies used by the data symbols, it is recommended to use frequencies between 500 Hz and 1.4 KHz. The former corresponds to a strip width of 66 to 100 pixels while the latter corresponds to a stripe width of 23 to 36 pixels.

2.2 Symbol duration

RS-FSK uses a symbol duration that equals to the receiving camera's frame duration. Since most of the cameras use a frame rate of 30 frame per second when capturing video, the default symbol duration is set to be 1/30 second. Note that the symbol can be configured by the use of optional field and PIB attributes in PPDU (see section 5.2).

2.3 Relation between signal frequency and the observed strip width in the captured image

In this section we describe the relationship between the transmitted frequency f and the strip width W. Although the standard usually does not specify how the receiver demodulate the transmitted frequency, this background knowledge is required to determine the set of frequencies that is used by the transmitters based on the specification of the receiving optical camera. Therefore, we choose to disclose the information here.

The strip width is defined as the number of pixels occupied by a set of bright strip (exposure during the transmitter is in the ON state) and a dark strip (exposure during the transmitter is in the OFF state) in a received image. Note that, for a square wave of frequency f Hz, the duration of a complete cycle is $1/f$ seconds. Therefore, for every $1/f$ seconds a camera exposes, it should be able to read out a pair of bright and dark strips in the received image. On the other hand, the time a camera spends to read out a row of pixels is denoted as its read-out duration T_r. Denote the width of a bright strip as W_b and the width of a dark strip as W_d, and the width of a pair of bright strip and dark strip as $W = (W_b + W_d)$. Therefore, in theory, the width of a pair of bright strip and dark strip can be found by $W = \frac{\frac{1}{f}}{T_r} = \frac{1}{fT_r}$ (see Figure AB). Note that W is a real number. In practice, a receiver would need to observe the width of a large number of pairs of strips to calculate the average number of rows occupied by a pair of strips, W', as an estimate of W, and demodulate the symbol by $f' = \frac{1}{W'T_r}$.

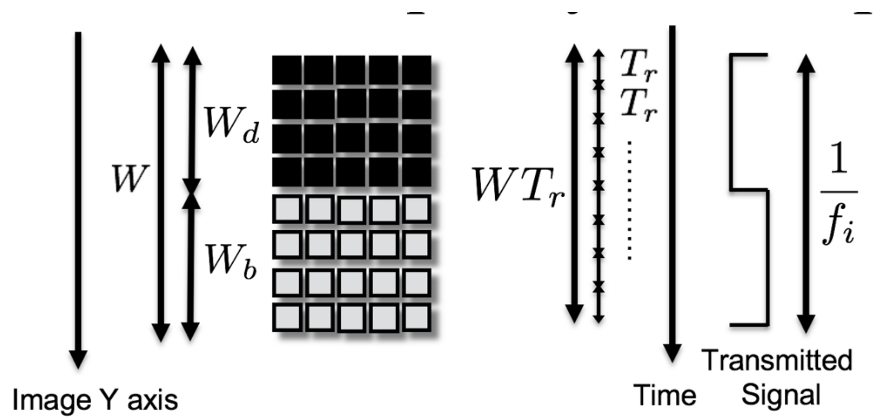

Figure AB – Illustration of the relation between transmitted RS-FSK signal frequency and the strip width in the corresponding camera captured image

2.4 Survey of parameters of commercially available image sensors

We have carried out experiments to survey the parameters of common commercially available image sensors, and the results are summarized in Table A.

Brand	Product Name	Image Resolution (X x Y)	Frame Rate (FPS)	Measured Read-out Duration (µs)	Gap between Frames (ms) / %
Apple	iPhone 6 plus	1920 x 1080	30	21.42	10.20 / 30.60%
	iPhone 5s	1920 x 1080	29.98	20.65	11.03 / 33.10%
	iPhone 4s	1920 x 1080	29.87	24.48	7.04 / 21.03%
HTC	New One	1920 x 1080	29.94	19.08	12.79 / 38.30%
Samsung	Galaxy S4	1920 x 1080	29.93	25.53	5.84 / 17.48%

Table A – Read-out duration of cameras

From the table, we can see that most cameras have a read-out duration of 20-30 microseconds. This provides motivation for the selection of frequency used by RS-FSK.

3.0 PHY Layer Dimming Method

RS-FSK supports dimming by changing the duty cycle of the transmitted signal. This allows the system to adjust the observed average brightness by human eyes. Note that the duty cycle is independent of and does not affect the transmitted signal frequency *f*, allowing the same demodulation scheme across different dimming settings (see Figure AC). However, as the duty cycle setting is configured to be further from 50% (i.e., very bright or very dark), the symbol error rate is expected to increase as in these cases it is more difficult to accurately determine the strip width from the captured image.

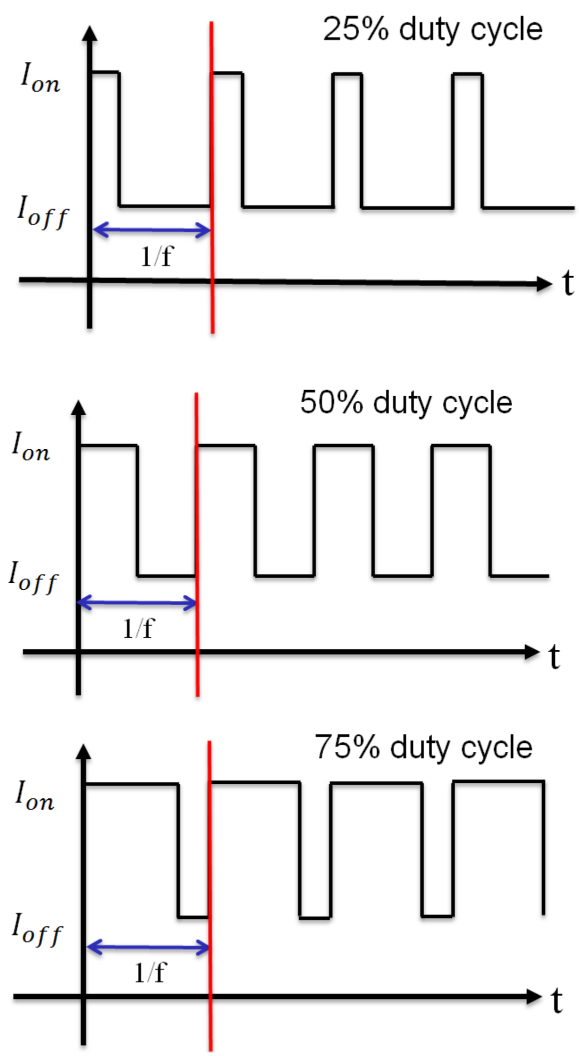

Figure AC – Examples of RS-FSK Dimming Method, with 25%, 50%, and 75% duty cycles

4.0 PPDU format

Due to convention, the left most field shall be transmitted or received first. All multiple octet fields shall be transmitted or received least significant octet first and each octet shall transmit or receive LSB first as well. The same transmission order should apply to data fields transferred between the PHY layer and the MAC sub-layer. The PPDU frame structure shall be formatted as Figure N.

Preamble (see 4.1)	Optional (see 4.2)	PSDU (see 4.3)	End Symbol (see 4.4)

Figure N – Format of the PPDU

4.1 Preamble field

The preamble field is used by the receiver to obtain optical clock synchronization with an incoming message. This will be the frequency baseline, which is denoted as *aPreambleFrequency* (check Section 5.1). The frequency used is 2.2 KHz, as specified in section 2.1. The duration of the preamble field is set to be one symbol duration, or 1/30 second (see section 2.2).

The preamble is a time domain sequence and does not have any channel coding or line coding. The same preamble shall be used for all PHY types. The number of repetitions of the preamble field can be extended by the MAC when requested for better synchronization.

4.2 Optional field

Optional field only appears during the synchronization section of the superframe (check section 6.0). Value of the optional field should be specific ratio of the preamble frequency, therefore, the frequency used in the PSDU cannot overlap with the one used in the optional field. If the optional field is assigned, then the data in the PSDU will be, and will only be, the parameter of the field of interest, no data payload can be assigned at the same time.

The optional field shall be formatted as shown in Figure Q. The optional fields in Figure Q (a) is used to indicate that the optional field is required, it should only present its appearance in the synchronization frame (check Section 6.2) of the superframe structure.

The optional field in Figure Q (b) shall be transmitted only when the transmitter informs the receiver about the frequencies used in the data symbols. The optional fields in Figure Q (c) shall be transmitted when one tries to configure the PIB values (check Section 5.0 for more details). Optional field in (b) and (c) shall never be used simultaneously since they correspond to different attributes.

Gap (see 4.2.1)	Frequency Labeling (see 4.2.2)	PIB Attribute (see 4.2.3)

(a) **(b)** **(c)**

Figure Q – PPDU optional fields

4.2.1 Gap field

Gap field is a blank field, which is used to indicate the start of the optional field. The frequency used by the gap field is defined as 10 times of the preamble frequency, i.e., 22 KHz. The duration of the field is one symbol duration.

4.2.2 Frequency Labeling field

Frequency labeling field is a value defined precisely at one and a half of the preamble frequency, i.e., 3.3 KHz. This is transmitted for one symbol duration. When this field is transmitted, it indicates the PSDU contains the data frequency that will be used in subsequent transmitted data frames (check Section 6.2), and the PSDU will operate without splitter symbol (check Figure M3). When transmitting frequency labels, the number of frequencies included in the PSDU should be a power of two. The frequency in the PSDU shall arranged sequentially by the transmitter in ascending order, starting from the frequency representing the smallest bit pattern (i.e., all '0' bit pattern) and ending from with the frequency representing the largest bit pattern (i.e., all '1' bit pattern). Consecutive frequency f_i and f_{i+1} should not remain the same. If f_i and f_j are the same, while $i \neq j$, the behavior is *undefined*.

4.2.3 PIB Attribute field

Directly assign the attribute frequency in this field to assign the value in PSDU field to the receiver, the format is predefined in Section 5.0.

4.3 PSDU field

If splitter symbol is required, the PSDU field has a variable length and carries the data of the PHY frame. The SS (Splitter Symbol) are introduced at the head and tail of each carried DS (Data Symbol) in Figure M. The head and tail SS are still appended if the PSDU has no payload (check Figure M2).

SS	DS 1	SS	DS 2	SS	…	SS	DS N	SS

Figure M – Format of the PSDU

SS	SS

Figure M2 – Format of the PSDU without data payload

SS is not necessary if the frame rate is matched between the transceiver and the receiver, this feature can be disabled through the optional field (check Section 4.2.3 for the format, Section 5.0 for all the attributes). In that case, the PSDU field has the format in Figure M3.

DS 1	DS 2	…	DS N

Figure M3 – Format of the PSDU without SS

Intuitively, when there is no payload, then the PSDU will cease to exist when SS is disabled.

4.4 End Symbol field

This field marks the end of the PPDU, simple but a necessity, in order for the receiver to acknowledge the end of this package. The end symbol frequency is defined as 0.75 times the preamble frequency, i.e., 1.65 KHz, and last for one symbol duration.

5.0 PHY constants and PIB attributes

This sub-clause specifies the constants and attributes required by the PHY.

5.1 PHY constants

The constants that define the characteristics of the PHY are presented in Table C. These constants are hardware and dependent and shall not be changed during operation.

Constant	Description	Value
aPreambleFrequency (aPF)	The preamble field frequency, this should be universal across different hardware that comply this definition.	2.2 KHz
aFrequencyLabelingRatio (aFLR)	This indicates the frequency ratio of the frequency labeling field ($aPF \times aFLR$).	1.5

Table C – PHY constants

5.2 PIB attributes

The PHY PIB comprises the attributes required to manage the PHY of a device. Each of these attributes can be assigned to the receiver through the optional field during the synchronization frame. The attributes contained in the PHY PIB are presented in Table D.

All the PHY PIB are labelled as frequency ratio, using *aPreambleFrequency* as the baseline. If one would like to extend the PIB attribute table, do not overlap the ratio with *aFrequencyLabelingRatio*, since it will seriously interfere the interpretation of all subsequent data symbols.

Attribute	Description	Frequency Ratio
*phyFrequencyLabeling**	Indicates the PSDU carries the frequencies that will be used for the data frame later on.	1.5
phyUseSplitterSymbol	Indicates the the PSDU carries the flag to toggle whether the device is going to use SSs or not.	0.83
phySplitterFrequency	Indicates the PSDU carries the splitter frequency. If the SS is already in used, it will use the original *phySplitterFrequency* until next cycle.	0.71
phySplitterDuration	Indicates the PSDU carries the duration of the SS. This is represented as a ratio of symbol duration to splitter duration in integer.	0.62
phySymbolDuration	Indicates the PSDU carries the duration of a data symbol in the PSDU. This is represented as a ratio of	0.55

	the symbol duration to 1/30 second in the base 2 log scale. For example, if the symbol duration is 1/120 second, then the PSDU would contain an integer -2. If the symbol duration is 1/15 second, then the PSDU would contain an integer 2. Note that this does not affect the duration of the preamble field and the optional field.	

Table D – PHY PIB attributes

* This PIB attribute is explicitly defined for future extensibility. For now, only *aFrequencyLabelingRatio* can control the frequency ratio of the Frequency Labeling field.

6.0 Superframe Structure

The transmitter needs to use a superframe structure to control the asynchronous operation between the transmitter and the receiver. A superframe is bounded by the transmission of a synchronization frame and a series of data frame. The receiver may enter a low-power mode if the synchronization frame indicates so, which is the inactive frame.

An example of a superframe structure is shown in Figure B. In this case, it includes the complete structure: synchronization frame (SF), data frame (DF) and inactive frame (IF).

SF (6.1)					DF (6.2)				IF (6.3)
FDP	SP 1	SP 2	…	SP N	DP 1	DP 2	…	DP N	

Figure B – An example of the superframe structure

6.1 Synchronization Frame (SF)

The synchronization frame contains all the instruction for the receiver to properly interpret the communication from the transmitter. It should start with a frequency definition packet (FDP), which enable the optional field Frequency Labeling field. If more configurations are required, append them after the FDP as the synchronization packets (SP).

6.2 Data Frame (DF)

Data frame contains all the packets with the actual data payloads.

6.3 Inactive Frame (IF)

The length of the inactive frame is defined by the *phyLowPower*. During this duration, receiver will not be able to receive data if if the transmitter is still transmitting. This is a logically defined frame. Transmitter isn't bound to this, i.e., the transmitter does not have to enter the low-power mode.

The duration of IF can be zero to maximized the transmission efficiency on the temporal scale. It is adjustable through the PHY PIB attribute.

7.0 MAC frame formats

This sub-clause specifies the format of the MAC frame (MPDU). Each MAC frame consists of the following basic components:

a) A MFH, which comprises frame control, sequence number and address information.
b) A MFDU, of variable length, which contains information specific to the frame type. Acknowledgement frames do not contain a payload.
c) A MFT, which contains a FCS.

The frames in the MAC sub-layer are described as a sequence of fields in a specific order. All frame formats in this sub-clause are depicted in the order in which they are transmitted by the PHY, from left to right, where the left most bit is transmitted first in time. Bits within each field are numbered from 0 (left most and least significant) to $k-1$ (right most and most significant), where the length of the field is k bits. Fields that are longer than a single octet are sent to the PHY in the order form the octet containing the lowest numbered bits to the octet containing the highest numbered bits.

For every MAC frame, all reserved bits shall be ignored upon receipt.

7.1 General MAC frame format

The MAC frame format is composed of a MAC frame header (MFH), a MAC frame data unint (MFDU), and a MAC frame tail (MFT). The fields of the MFH appear in a fixed order; however, the addressing fields may not be included in all frames. The general MAC frame shall be formatted as illustrated in Figure Y.

Octets: 1	1	0/2	0/2	variable	2
Frame Control	Sequence Number	Destination PAN Address	Source PAN Address	Frame Payload	FCS
		Addressing fields			
MFH				MFDU	MFT

Figure Y – General MAC frame format

7.1.1 Frame Control field

The Frame Control field is 1 octet in length and contains information defining the frame type, addressing fields, and other control flags. The frame control field shall be formatted as illustrated in Figure J. Reserved bits are set to zero on transmission and ignored on reception.

Bits: 0-2	3	4	5	6	7

Frame Type	Reserved	Security Enabled	Frame Pending	Destination Addressing Mode	Source Addressing Mode

Figure J – Format of the frame control field

7.1.1.1 Frame Type subfield
The Frame Type subfield shall be set to one of the non-reserved values listed in Table Q.

Frame Type value b_2 b_1 b_0	Description
000	Reserved
001	Data
010	Command
011	Security Configuring Mode
100-111	Reserved

Table Q – Values of the Frame Type subfield

7.1.1.2 Security Enabled subfield
The Security Enabled subfield is 1 bit in length, and it shall be set to one if the frame is protected by the MAC sublayer and shall be set to zero otherwise. Prior to enable this field, the transmitter should configure the receiver into Security Enabled prepared state through the Frame Type subfield.

7.1.1.3 Frame Pending subfield
The Frame Pending subfield is 1 bit in length and shall be set to one if the device sending the frame has more data for the recipients. This subfield shall be set to zero otherwise.
The Frame Pending subfield shall be used only during the DF, at SF it shall be set to zero on transmission and ignored on reception.

7.1.1.4 Destination Addressing Mode subfield
If this subfield is equal to zero, the Destination PAN Address shall not be included.

7.1.1.5 Source Addressing Mode subfield
If this subfield is equal to zero, the Source PAN Address shall not be included.

7.1.2 Sequence Number field
The Sequence Number field is 1 octet in length and specifies the sequence identifier for the frame.

7.1.3.1 Destination PAN Address field

The Destination PAN Address, when present, is 2 octets in length, and specifies the address of the intended recipient of the frame. A 16-bit value of 0xFFFF in this field shall represent the broadcast address, which shall be accepted as a valid 16-bit address by all devices currently listening to the channel.

This field shall be included in the MAC frame only if the Destination Addressing Mode subfield of the frame control field is nonzero.

7.1.3.2 Source PAN Address field

The Source PAN Address, when present, is 2 octets in length, and specifies the address of the originator of the frame. This field shall be included in the MAC frame only if the Source Addressing Mode subfield of the frame control field is nonzero.

7.2 MFDU

The MFDU contains the frame payload, which has a variable length and contains information specific to individual frame types. If the frame control is configured to Security Enabled previously, then the frame payload is protected as defined by the security suite selected at that time.

7.3 MFT

Currently MFT contains only the frame checksum (FCS). The FCS field is 2 octets in length and is explained in somewhere else in the document. The FCS is calculated over the MFH and MSDU part of the frame. The FCS shall be only generated for payloads greater than zero bytes.

8.0 MAC constants and PIB attributes

This sub-clause specifies the constants and attributes required by the MAC sublayer.

8.1 MAC constants

8.2 MAC PIB attributes

The MAC PIB comprises the attributes required to manage the MAC sublayer of a device. The attributes contained in the MAC PIB are presented in Table G.

Attribute	Description	Default
macAttribute	Attribute description.	Default value.

Table G – MAC PIB attributes