Flexilink

A New Low-Latency Network for AV

A Unified Low Latency Network Architecture for Multichannel Live and Interactive Media

John Grant, Kate Grant
Nine Tiles, Cambridge, UK

Yonghao Wang, Jeremy Foss
School of Digital Media Technology, Birmingham City University, UK

The Problem with Traditional IP Media

Media transfer over an IP network suffers end-to-end latency, specific cases being in the main environments the signal encounters:

In a media processing system:
- analogue / digital / analogue conversion
- networking and routing
- digital consoles
- computer systems (e.g. with software plugins)

In IP networks - Buffering, especially at Internet nodes (also including intranets)

Current QoS Solutions - traffic engineering and over-provisioning - complicate the system, increase power requirements and add cost

Professional Audio Networking Industry – are reluctant to use current Internet solutions for time critical applications. Live audio applications typically use closed audio specific networking technologies with modified layer-2 or layer-3 protocols to utilise the lower network layer. These applications typically use AES50, EtherSound, CobraNet etc.

Interactive TV – Next generation interactivity will allow viewers to interact with on-screen and second screen apps, including viewer-to-viewer interaction. This will require a highly responsive support of video traffic on carrier networks.

Summary It is difficult to achieve an architecture which supports time critical data simultaneously with best effort data. Connectionless packet architectures are inevitably a problem for deterministic low latency data. There are also issues for multichannel streams with a range of sampling frequencies and variable bit lengths, and the need for flexible routing and channel assignments. Current multiplexing methods are insufficient to support them without sampling rate conversion and data format rectification.

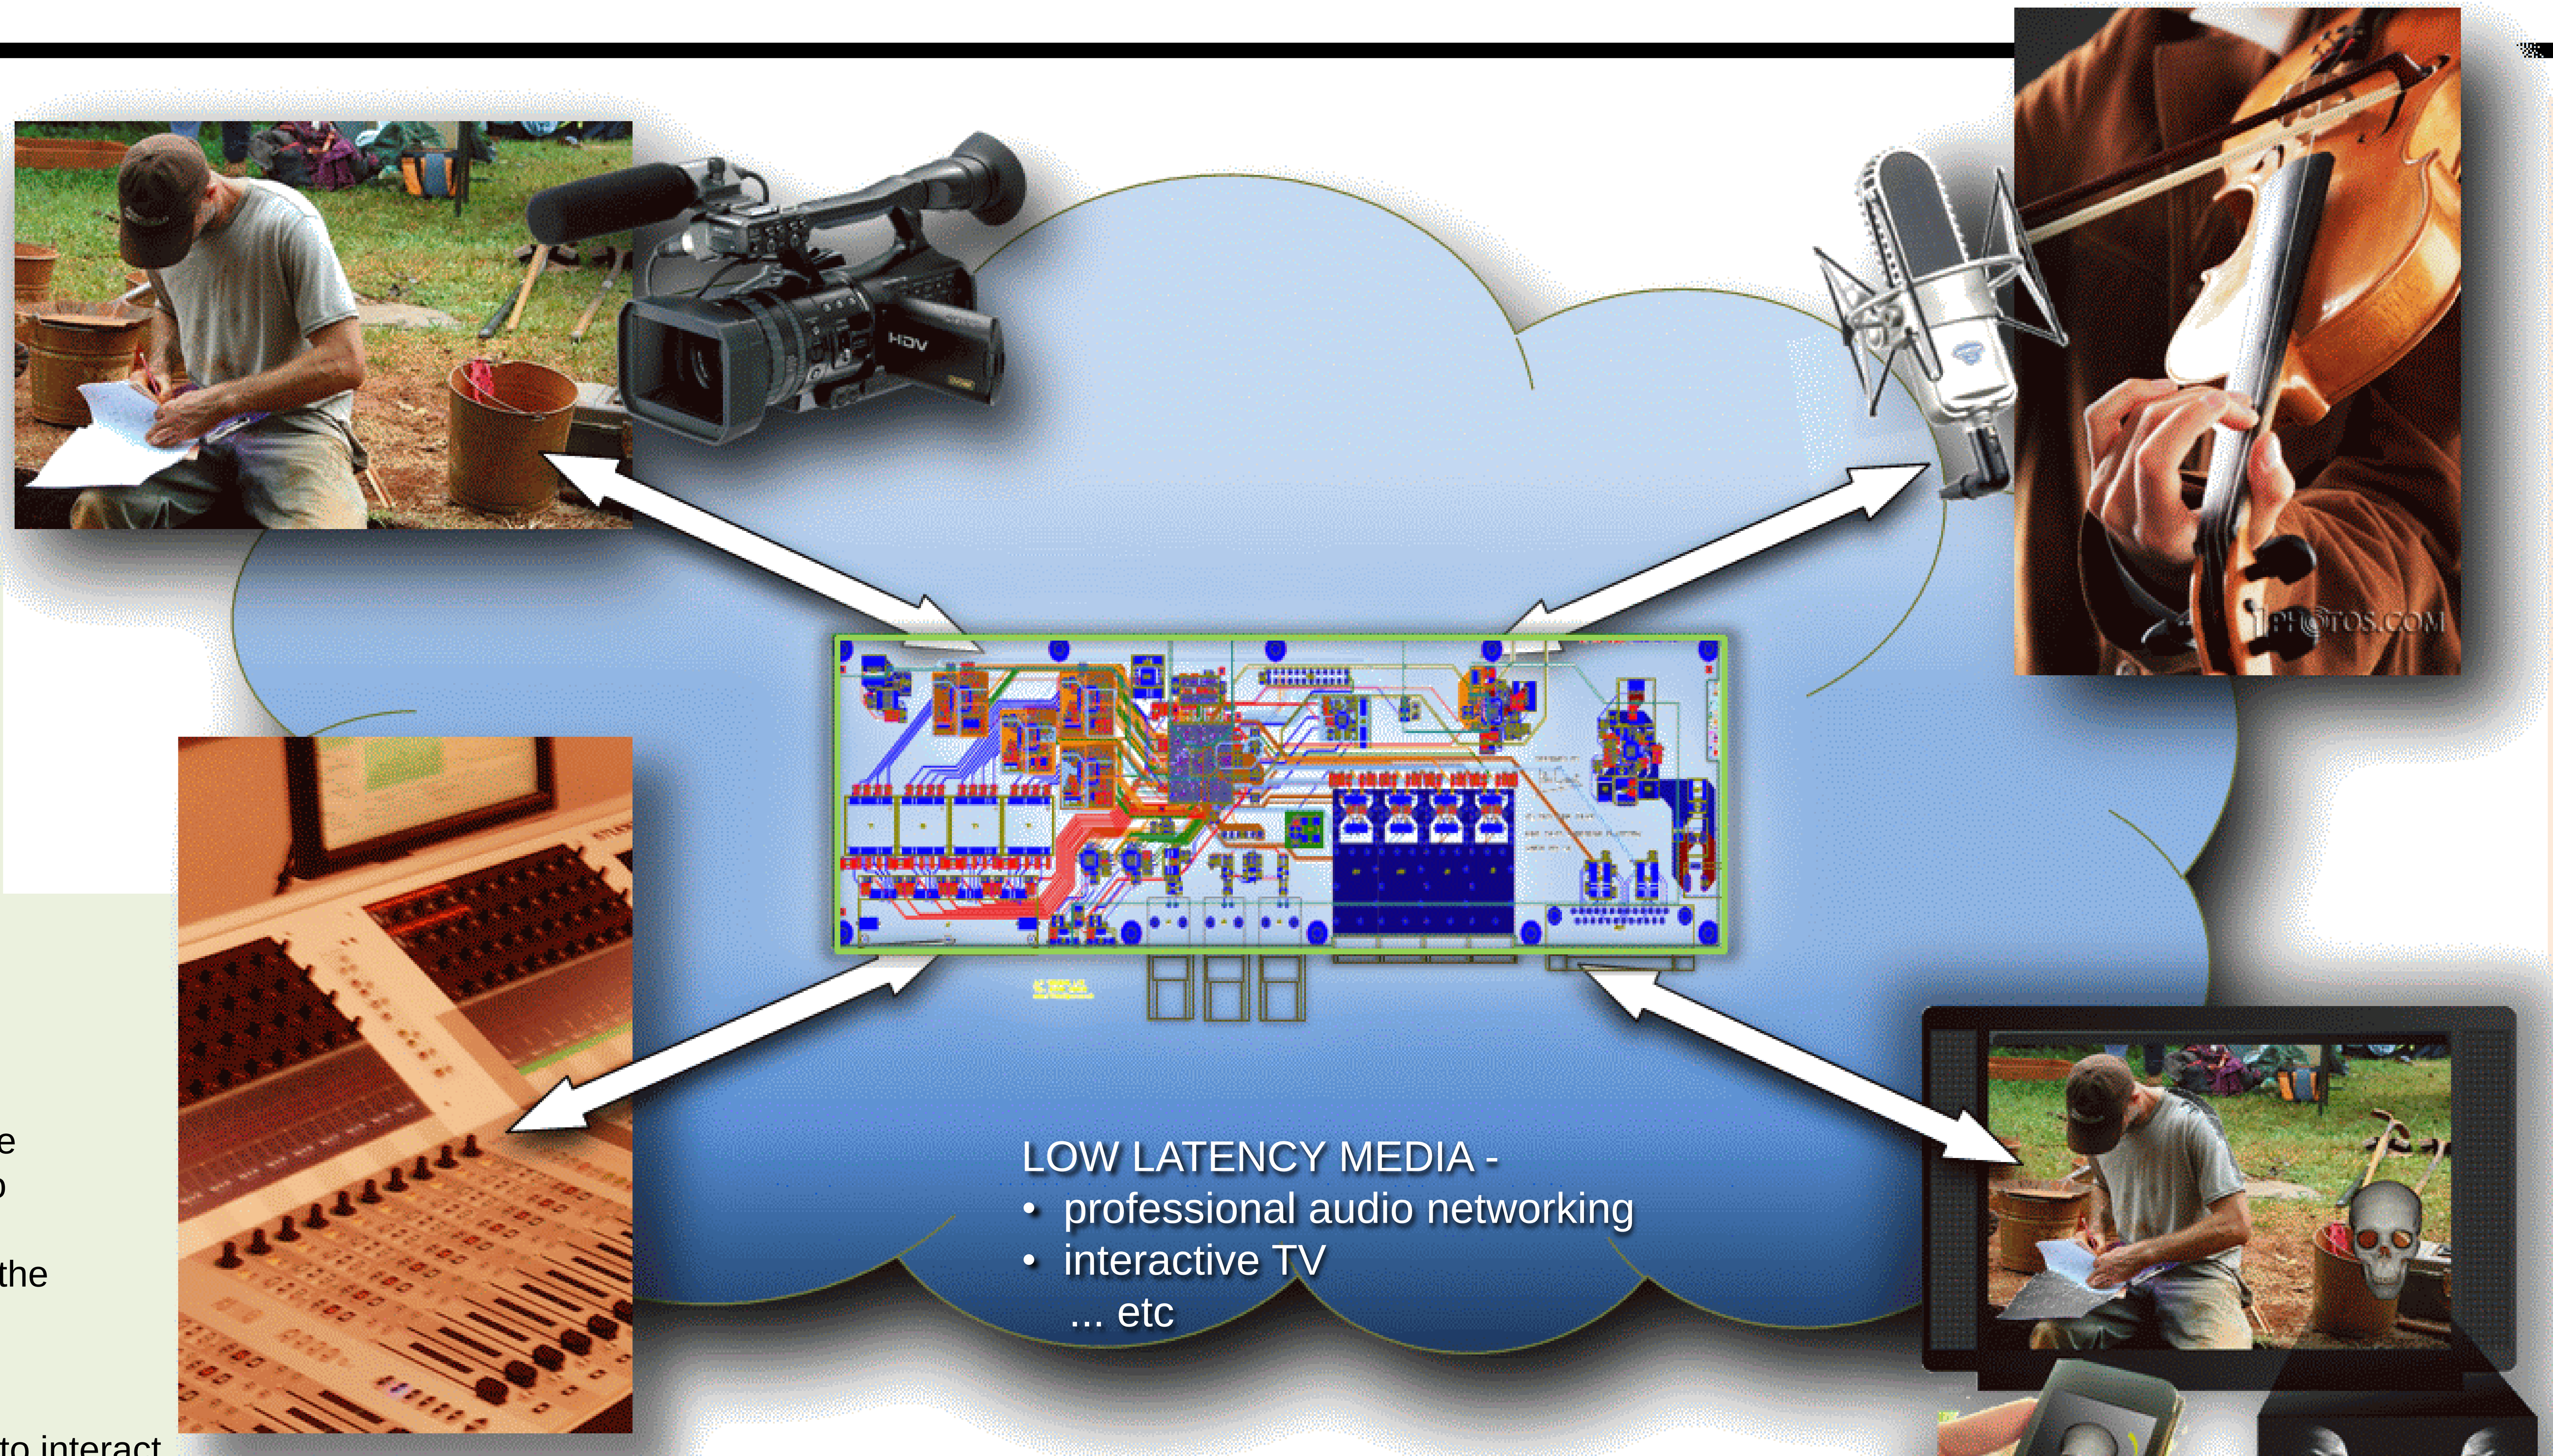

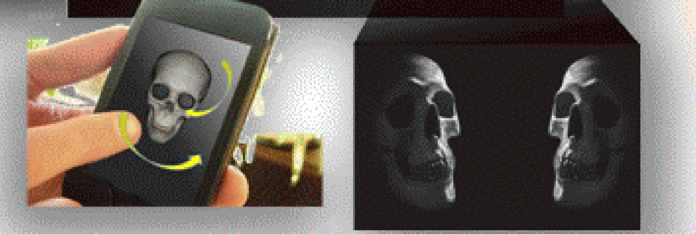

FLEXILINK SOLUTION

Combines the advantages of TDM and *best effort* networks
- Supports best effort data and time deterministic data (such as audio sample packets).
- Interworks compatibly with existing network infrastructure and protocols
- The Physical Layer is a time deterministic bit pipe, so a time deterministic logical control is implemented

Allows a guaranteed Quality of Service (QoS) and expected Quality of Experience (QoE) for the higher layer protocols
- Exact time delay / transmission of time-critical data can therefore be estimated

Layer 2 protocols defined in FLEXILINK
- Developed along with a prototyped network processor architecture and interface cards
- Compatibility with existing Ethernet infrastructure
- Operates at full-duplex mode (where non-deterministic CSMA-CD can be avoided)

FLEXILINK OPERATION

Assume three data categories:
- **Synchronous Flow (SF) packets** for audio/video and other time deterministic data.
- **Asynchronous Flow (AF) packets** for best effort data.
- **Control Message (CM)** for session control and link management.

Synchronous Flow (SF)
Time-deterministic media packets:

Asynchronous Flow (AF)
Best effort data packets:

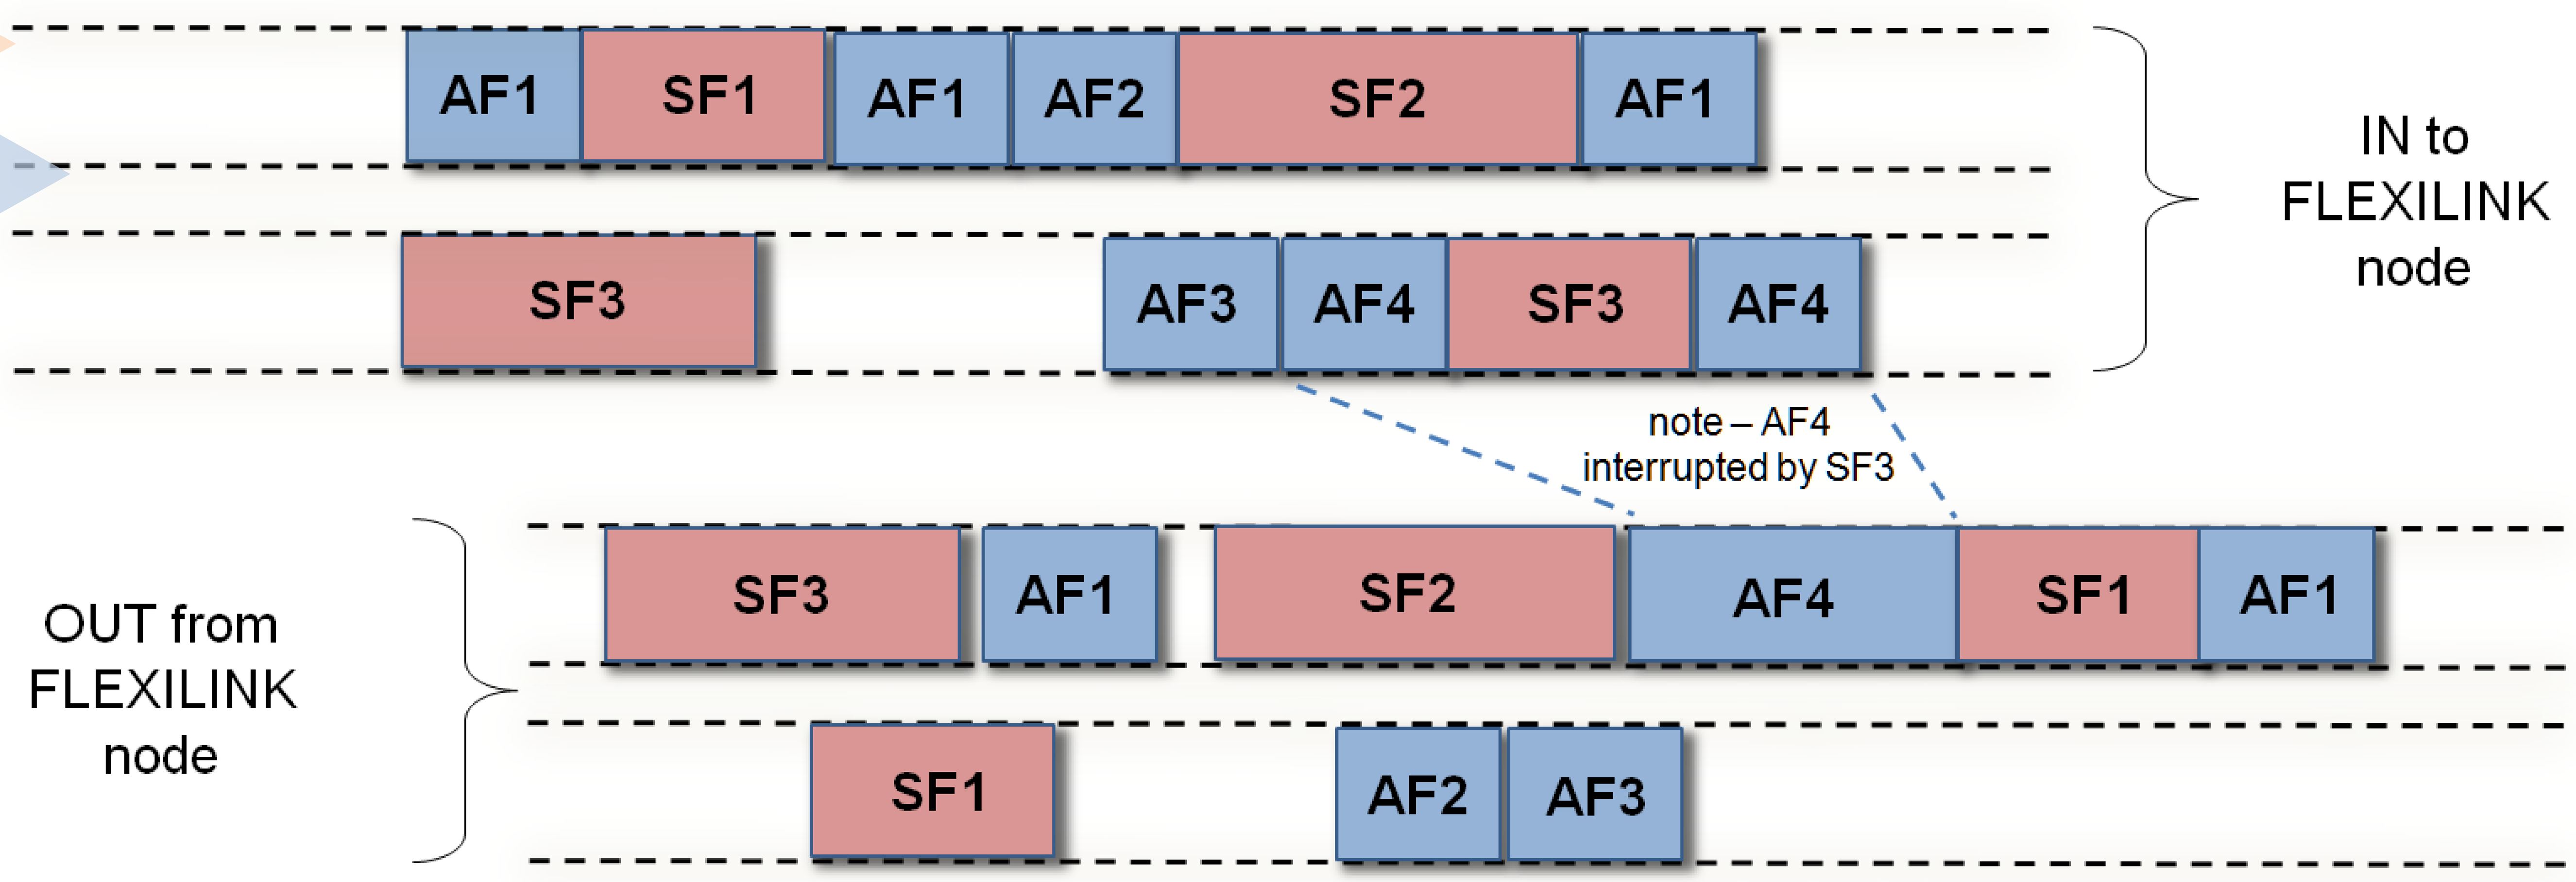

Here, three SFs are shown here with SF1 being multicast. Resources (i.e. bandwidth requirements) need to be reserved when the link is established to support deterministic media data. SF data packets are identified by their position in the stream (*c.f.* time slots in a TDM frame).

AFs are transmitted in available time in the frame not allocated to SF data. Since SF packets are likely to vary in length, AF allowance will also be of variable length

John Grant, Kate Grant
Nine Tiles, Cambridge, UK
j@ninetiles.com
kate@ninetiles.com

Jeremy Foss, Yonghao Wang
Birmingham City University, UK
jeremy.foss@bcu.ac.uk
yonghao.wang@bcu.ac.uk

Nine Tiles

BIRMINGHAM CITY University

Flexilink

A New Low-Latency Network for AV

A Unified Low Latency Network Architecture for Multichannel Live and Interactive Media

John Grant, Kate Grant
Nine Tiles, Cambridge, UK

Yonghao Wang, Jeremy Foss
School of Digital Media Technology, Birmingham City University, UK

FLEXILINK Architecture In The Network Node

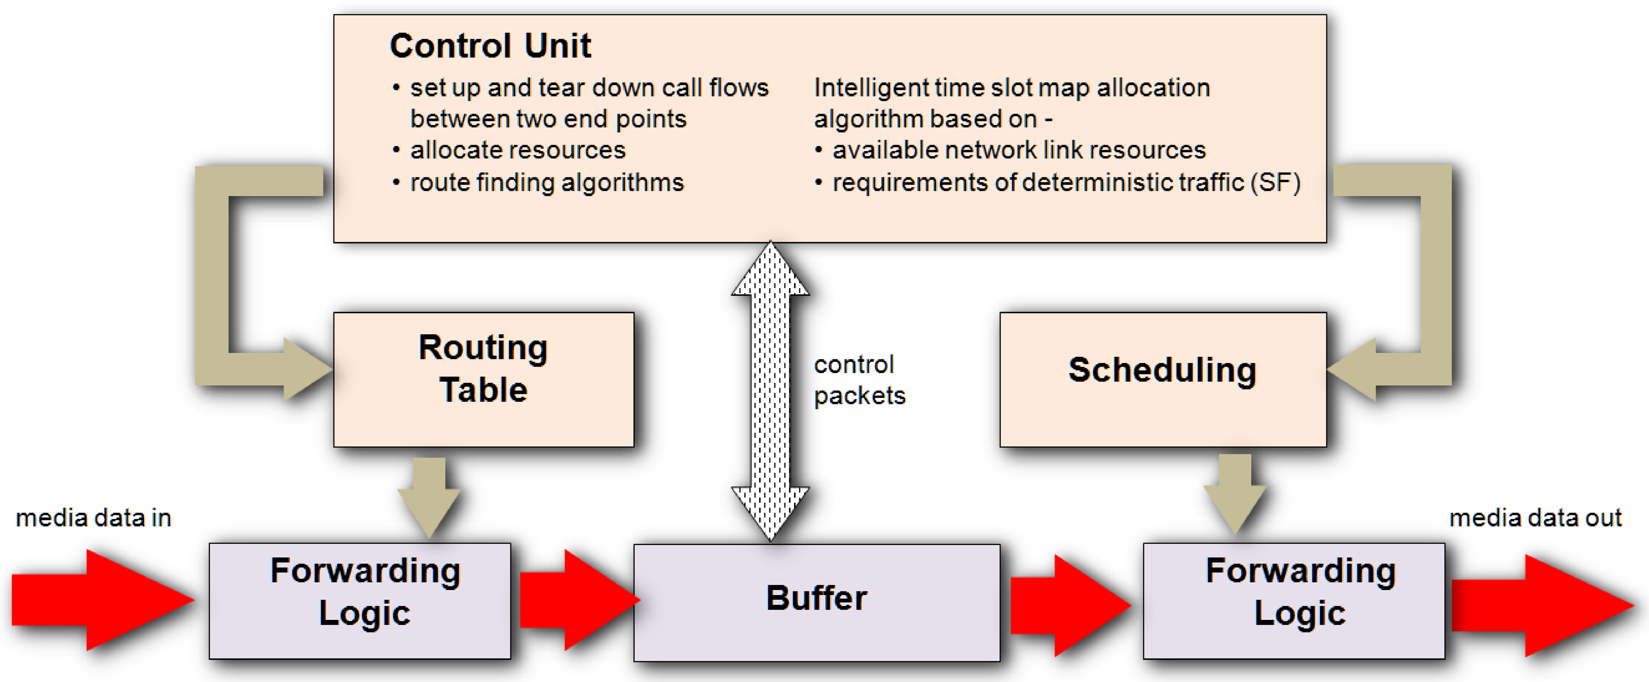

Simplified Packet Headers

Headers contain only packet length information (size of audio data in this packet); minimises header cost; length of header is variable to minimise overhead

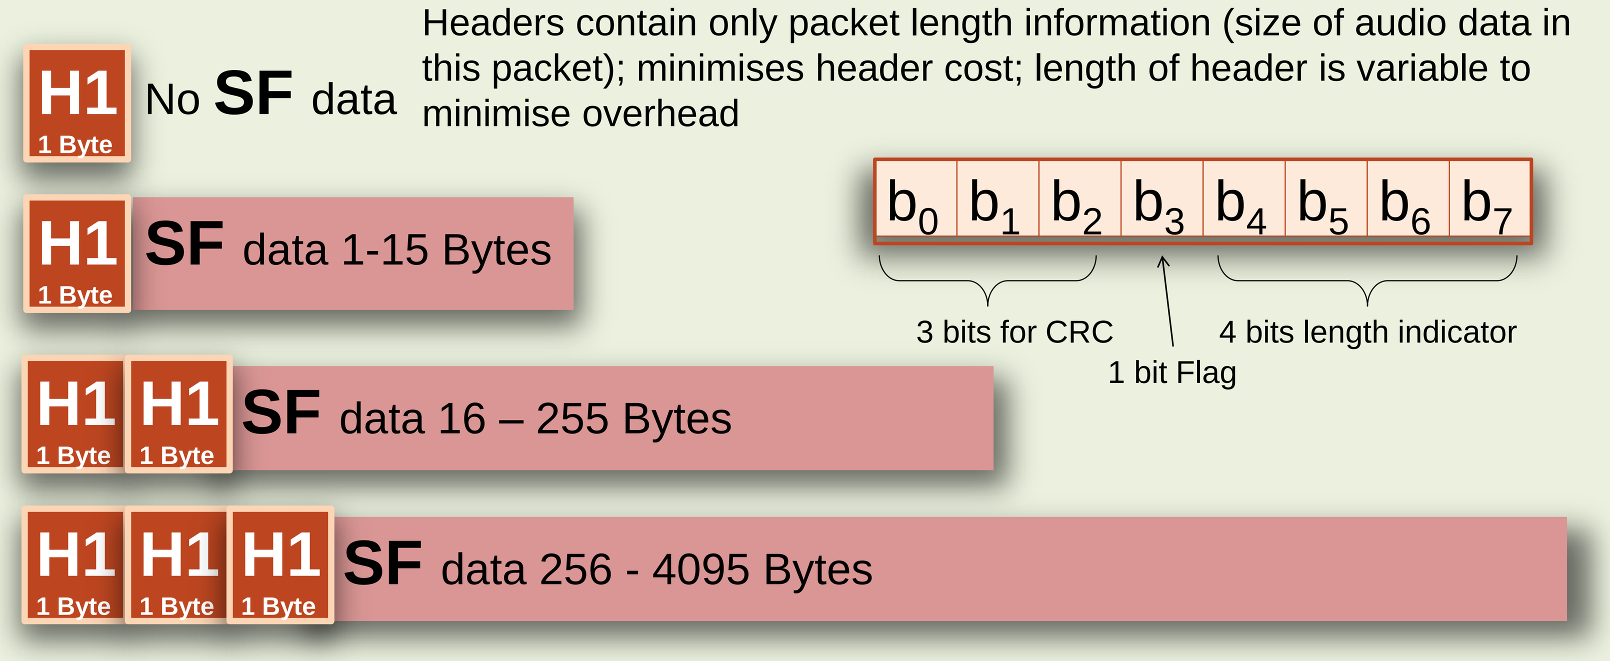

Interface architecture

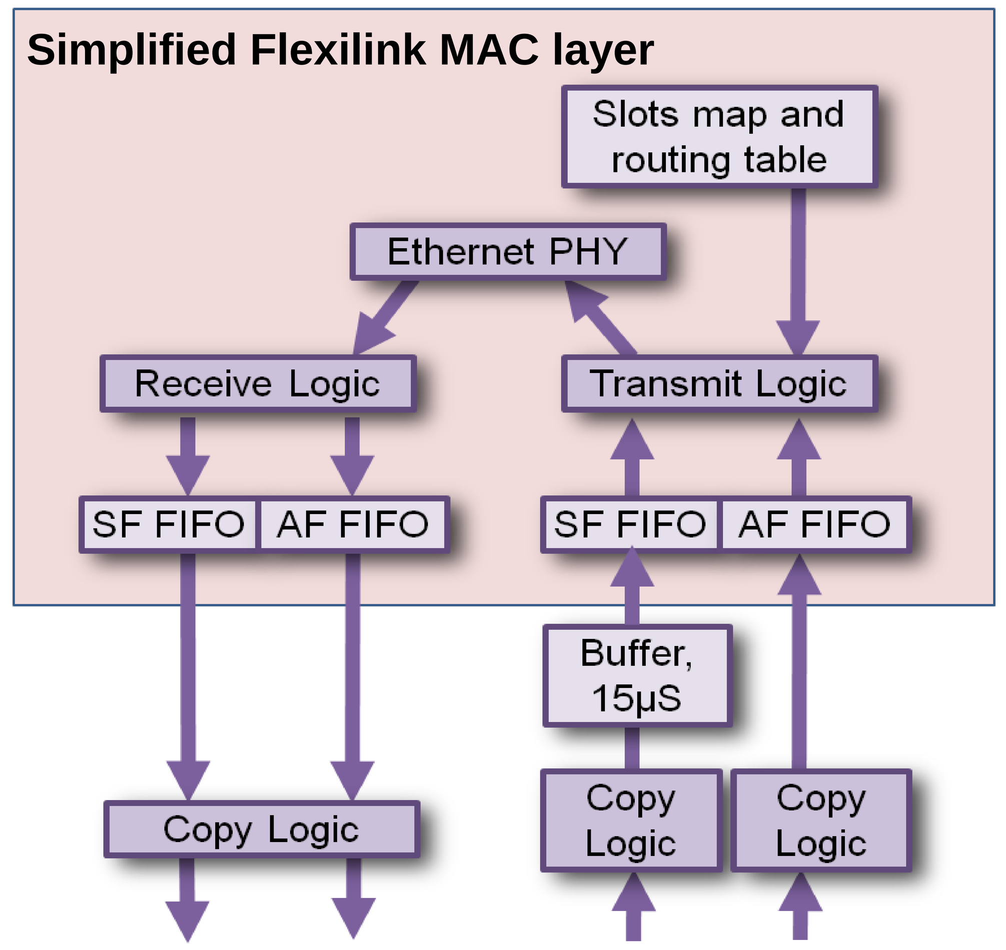

Interface architecture to support Flexilink

- Ideally, for maximum compatibility, Flexilink should be able to use the existing physical network interface. However to support the proposed Flexilink protocol, a new media access control (MAC) layer architecture needs to be considered, which allows AF and SF to be treated differently.
- AF and SF have separate buffers and copy logic allocated for them. The SF input has a 15µs buffer from which the packets are taken to fill the output slots
- Multicasting can easily be achieved simply by copying the same packet to several different outputs
- The SF buffer, containing the last 15µs of data from each input, is small enough to go in internal block RAM in the FPGA.
- The larger AF buffer is in external dRAM, and contains one queue per output.

Support In Standard Ethernet Frame

SF and AF access may be implemented over an existing point-to-point link mechanism in order to utilise current network infrastructure. A practical implementation of the layered traffic model is shown here. The Frame layer can be a standard fixed size Ethernet frame, so the position of an SF packet in relation to the start symbol of frame can be used as a reference for identification of the SF data packets (i.e. t_1 for SF1, t_2 for SF2, etc.) Consequently labels for SF identification are not required.

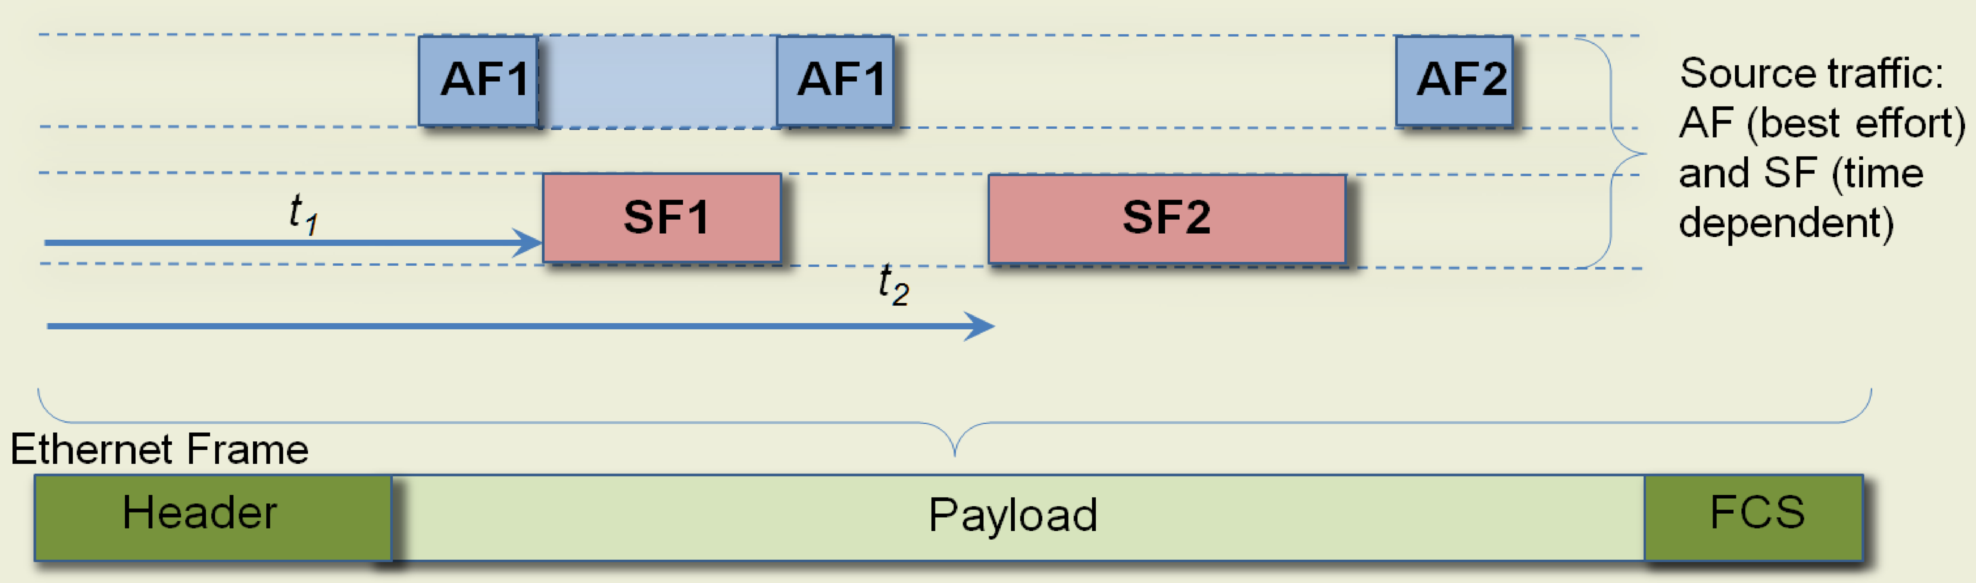

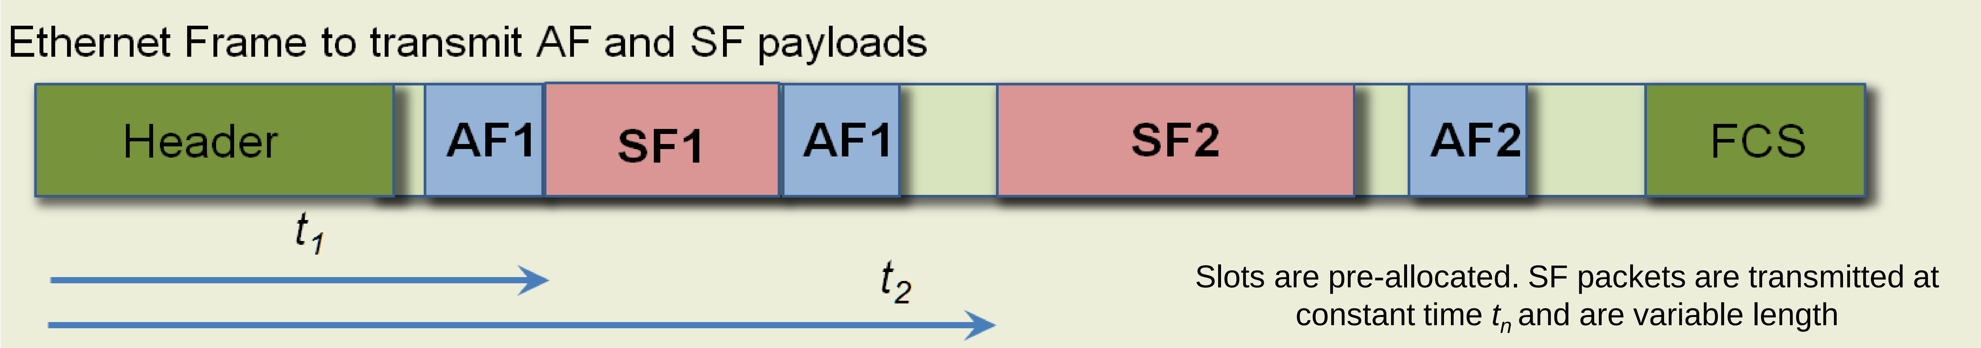

Supporting Flexible Multichannel Audio Streams With Different Sampling Frequencies

The current problem in networking audio traffic is the lack of flexibility to support arbitrary numbers of audio channels with different sampling frequencies and of compressed or uncompressed data format. In FLEXILINK multiple channels of different sampling frequencies can be readily supported (as long as the link capacity is sufficient for the SFs).

For Ethernet physical media the Ethernet jumbo frame format maximises the capacity of available bandwidth and minimise the cost of interframe gaps.

In **FLEXILINK**:

Time slot map

- allocated within the payloads of the frames
- allocation algorithm will ensure the SF streams are distributed evenly

Phase control algorithm – to ensure that each output slot always occurs after the input slot from which its packet comes, but no longer than 15µs after.

Precision Time Protocol can be adopted for accurate timing and provision of clocking.

The Implementation of FLEXILINK on Ethernet

Ethernet Frame structure for Flexilink adopts the Jumbo Frame format (i.e. for payload size > 1500 bytes for common Ethernet frame)

However - some Ethernet PHY chips do not support frames larger than 9000 B. To accurately access the allocation slots and the positions for the SF packets, it is preferred to have a fixed allocation period at the frame layer. In this case the allocation period of 124.96µS is chosen (i.e. slightly faster than 125 µS used in SDH, Firewire and full speed USB, etc).

The network devices can use standard AES51 negotiation packets (standard Ethernet MAC packets) to set up the links. Once both ends are in Flexilink mode, a "Reduced Jumbo Frame" (RJF) format is utilised to maximise the payload. RJF eliminates some unnecessary parts of the standard Ethernet frame (for example, the source and destination addresses, because all frames are sent from the node at one end of the link to the node at the other) to give more payload space to the allocation periods.

RJF Format:

- 2 bytes preamble + Start Frame Delimiter (SFD).
- 5 bytes AES51 packet type and timing information
- 7785 bytes payload data
- 4 bytes FCS
- 14 bytes interframe gap (IFG)

Total frame size including IFG is 7810 bytes

Two successive RJFs combine together to make a 124.96µS allocation period (AP) at full duplex 1 Gigabit Ethernet link.

This design is to guarantee that 8000 allocation periods/second can be transmitted over 1 Gigabit Ethernet link. Each allocation period has 15570 bytes payload space to transmit SF and AF traffic. The theoretical bandwidth utilisation can be up to 99.6%. The FCS is only used to check that the link is working reliably.

Routing of AFs and SFs does not wait until the FCS has been received, and for many media formats it is better to deliver data with a few bit errors than to discard whole frames.

Related Publication

Flexilink: A unified low latency network architecture for multichannel live audio; Wang, Grant, Foss; 133rd AES Convention, San Francisco USA, October 2012