

Spike Processing on an Embedded Multi-task Computer: Image Reconstruction

C. Luján-Martínez, A. Linares-Barranco, M. Rivas-Pérez,
Á. Jiménez-Fernández, G. Jiménez-Moreno and A. Civit-Balcells

Dept. de Arquitectura y Tecnología de Computadores,
Universidad de Sevilla, Sevilla, Spain.
{cdlujan,alinares,mrivas,ajimenez,gaji,civit}@atc.us.es

Fifth Workshop on Intelligent Solutions in Embedded Systems.

WISES 07, U. Carlos III, Madrid, Spain.
June 2.007.



Address-Event-Representation: AER

I. Neuronal Networks

Neuro-inspired systems

are electronic circuits that mimic the processing of the biological nervous system.

In biology,

the neural system is divided in layers of neurons. These layers are processing the information from the sensors.


- Each neuron of a layer is connected to a projective field of neurons in other layers, following a weighted pattern.
- Neurons communicate each others using ionic pulses, called *spikes*, with time constants of milliseconds.
- Information resides in the pulses frequency.

Address-Event-Representation: AER

II. Virtual Neuronal Point-to-point Communication

Problem:

The connectivity between layers implies a huge number of point to point connections and chips have limited the number of pins.



Solution:


Temporal multiplexation of neuron's pulses in a high speed digital bus, Address-Event-Representation.

Address-Event-Representation: AER

II. Virtual Neuronal Point-to-point Communication

Problem:

The connectivity between layers implies a huge number of point to point connections and chips have limited the number of pins.



Solution:

Temporal multiplexation of neuron's pulses in a high speed digital bus, Address-Event-Representation.

Address-Event-Representation: AER

III. AER Fundamentals

Each time a cell on a sender device generates a spike,

a digital word representing a code or address for that cell is placed on the external inter-chip digital bus, the AER bus. This word is called *event*.

In the receiver chip,

the spikes or events are guided to the cells whose code or address appeared on the bus.

In this way,

cells with the same address in the emitter and receiver chips are virtually connected by streams of spikes.



Address-Event-Representation: AER

IV. More AER Fundamentals and Some Low Level Features

Additional handshaking lines,

Acknowledge and Request, are used for completing the asynchronous communication.

In addition,

transmitting the cell addresses allows performing extra operations on the events while they travel from one chip to another. The output of a silicon retina

- can be easily translated,
- scaled, or
- rotated by simple mapping operations on the emitted addresses,

making AER not only a communication channel.



Address-Event-Representation: AER

V. A Vision Scheme: Rate-Coded AER

In artificial vision systems based in AER,

it is widely used the *rate-coded* AER

In this scheme,

each cell corresponds to a pixel and its activity is transformed into pixel event frequency.

Although,

it is inefficient for conventional image transmission:

Monochrome VGA resolution yields a peak rate of

$$(480 \times 640 \text{ pixels/frame}) \times (256 \text{ spikes/pixel}) \times (25 \text{ frames/s}) \times (19 \text{ bit/spike}) = 37 \text{ Gbit/s.}$$

On the other hand,

the lost of some events does not mean a degradation in the application.

Address-Event-Representation: AER

VI. Rate-coded AER: Reducing the Traffic

Transmitting preprocessed images instead of raw images,

such as edges or contrast, in which 20 gray levels are satisfactory and only (1-10%) of pixels will present appreciable contrast.

2~3 orders of magnitude

In addition,



present day AER hardware uses image resolutions between 64×64 and 128×128 pixels at the most.

1~2 orders of magnitude



Address-Event-Representation: AER

VII. AER Multistage Systems



In multistage systems,

- events are generated at the front end, they travel and are processed down the whole chain (without waiting to finish processing each frame). Also,
- information is reduced after each stage, thus reducing the event traffic.

AER-tools and Event Rates:

- Rome PCI-AER → 1 Mevents/s
- CAVIAR PCI-AER → 8 Mevents/s
- USB-AER → 25 Mevents/s
- USB2AER → 5 Mevents/s
- mini USB-AER → 300 Kevents/s

They

- are mainly based on reconfigurable hardware.
- achieve a high bandwidth.
- useful for mapping, monitoring and sequencing.
- But a PC is needed for the high level spike processing.



AER-tools and Event Rates:

- Rome PCI-AER → 1 Mevents/s
- CAVIAR PCI-AER → 8 Mevents/s
- USB-AER → 25 Mevents/s
- USB2AER → 5 Mevents/s
- mini USB-AER → 300 Kevents/s


They

- are mainly based on reconfigurable hardware.
- achieve a high bandwidth.
- useful for mapping, monitoring and sequencing.
- But a PC is needed for the high level spike processing.




The Embedded Computer


Connex board:



Wifistix board:



Tweener board:



Hardware:

- Intel XScale PXA255 400 MHz (32 bit processor).
- 32 KB data cache and 32 KB instruction cache
- 64 MB RAM (with MMU)
- 84 GPIOs (serial ports, I2C, PWM, LCD, USB client 1.1,...)
- 16 MB of Flash mem.
- wireless connectivity (Bluetooth or IEEE 802.11b).

Software:

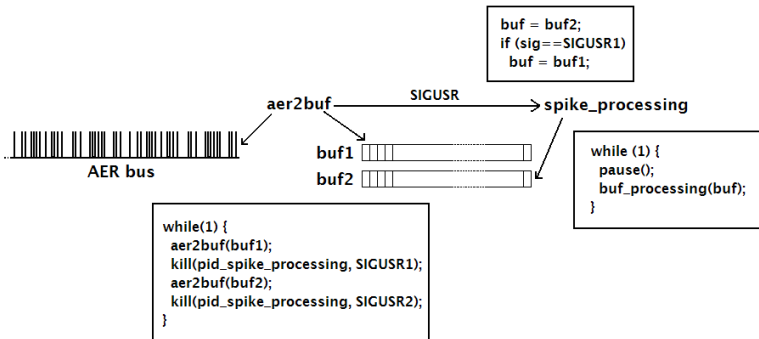
- GNU/Linux 2.6 kernel.
- Buildroot.
- uClibc.



Spike Processing over Multi-task


I. Software Architecture

- for high level spiking processing
- over a multi-task system
- when receiving events from the AER bus.



Spike Processing over Multi-task

II. Execution Flowchart Possibilities



Filling the buffer lasts more than consuming it, AER task B.

An acknowledge signal is used when AER task B has finished using the buffer.

AER task B aborts and starts using an updated buffer.



Image Reconstruction from Event Stream

I. From asynchronous AER to synchronous frame based representation video

- If T_{frame} is the duration of a single frame, \Rightarrow
 $t = n \times T_{frame}$ ($n \in [0, \infty)$), called the integration time.




Then,

- 1 for each event address, $[(x, y)] \leftarrow [(x, y)] + 1$.
- 2 at $t = (n + 1) \times T_{frame}$, the content of the 2-D memory is reset.



Image Reconstruction from Event Stream

II. Software Architecture and Scenario



Results

I. Processes vs Threads & Timer Interrupts

Event Rate, ER, and Worst Event Rate, WER, in Kevents/s:

| Test | Timer | No other process | | Other processes | |
|-----------|---------|------------------|-----|-----------------|-----|
| | | ER | WER | ER | WER |
| Processes | 100 Hz | 540 | 450 | 530 | 200 |
| Threads | 100 Hz | 770 | 620 | 770 | 259 |
| Processes | 1000 Hz | 500 | 430 | 500 | 430 |
| Threads | 1000 Hz | 775 | 660 | 770 | 660 |

By setting HZ to 1000

the ER is enclosed to a “controlled” interval.



Results

I. Processes vs Threads & Timer Interrupts

Event Rate, ER, and Worst Event Rate, WER, in Kevents/s:

| Test | Timer | No other process | | Other processes | |
|-----------|---------|------------------|-----|-----------------|-----|
| | | ER | WER | ER | WER |
| Processes | 100 Hz | 540 | 450 | 530 | 200 |
| Threads | 100 Hz | 770 | 620 | 770 | 259 |
| Processes | 1000 Hz | 500 | 430 | 500 | 430 |
| Threads | 1000 Hz | 775 | 660 | 770 | 660 |

By setting HZ to 1000

the ER is enclosed to a “controlled” interval.



Results

II. Scheduling Policies

GNU/Linux 2.6:

SCHED_RR:

each process is only allowed to run for a maximum time quantum.

SCHED_FIFO:

a simple scheduling algorithm without time slicing.

SCHED_OTHER:

dynamic priority time-sharing scheduler, based on the nice level and increased for each time quantum the process is ready to run, but denied to run by the scheduler.

- “scheduling” $\in O(1)$ + pre-emptive code \Rightarrow better response time.

Ruled by SCHED_RR

“threads implementation” remains **840 KEvents/s.**



Results

II. Scheduling Policies

GNU/Linux 2.6:

SCHED_RR:

each process is only allowed to run for a maximum time quantum.

SCHED_FIFO:

a simple scheduling algorithm without time slicing.

SCHED_OTHER:

dynamic priority time-sharing scheduler, based on the nice level and increased for each time quantum the process is ready to run, but denied to run by the scheduler.

- “scheduling” $\in O(1)$ + pre-emptive code \Rightarrow better response time.

Ruled by SCHED_RR

“threads implementation” remains **840 KEvents/s.**



Results

III. AER Communication vs Spike Processing

The minimum time between events, the system is able to achieve, is

1.16 μs .

So, the mean Inter-Spike-Interval, ISI, for

SCHED_OTHER:

ISI = **1.29 μs**

SCHED_RR:

ISI = **1.19 μs**

Therefore, the threads implementation presents a deviation from the best of

SCHED_OTHER:

11%

SCHED_RR:

2.5%



Results

III. AER Communication vs Spike Processing

The minimum time between events, the system is able to achieve, is

1.16 μ s.

So, the mean Inter-Spike-Interval, ISI, for

SCHED_OTHER:

ISI = **1.29 μ s**

SCHED_RR:

ISI = **1.19 μ s**

Therefore, the threads implementation presents a deviation from the best of

SCHED_OTHER:

11%

SCHED_RR:

2.5%




Results

IV. Real Applications

CAVIAR: a neuromorphic vision system totally based on AER where

the maximum throughput ER takes place at the output of the silicon retina for real applications.



$$\text{ER} \in [8, 150 \text{ Kevents/s}] \Rightarrow \text{ISI} \in [6.66 \mu\text{s}, 123.8 \mu\text{s}]$$

[P. Lichtsteiner and T. Delbruck. 64×64 Event-Driven Logarithmic Temporal Derivative Silicon Retina. *Proceedings of IEEE Workshop on Charge-Coupled Devices and Advanced Image Sensors*, pages 157-160, 2005.]

Wich means

from $5.5 \mu\text{s}$ to $122.6 \mu\text{s}$, or from 2200 to 50000 instructions, on a 32-bit μp at 400 MHz for any kind of high level spike processing.




Results

IV. Real Applications

CAVIAR: a neuromorphic vision system totally based on AER where

the maximum throughput ER takes place at the output of the silicon retina for real applications.



$$\text{ER} \in [8, 150 \text{ Kevents/s}] \Rightarrow \text{ISI} \in [6.66 \mu\text{s}, 123.8 \mu\text{s}]$$

[P. Lichtsteiner and T. Delbruck. 64×64 Event-Driven Logarithmic Temporal Derivative Silicon Retina. *Proceedings of IEEE Workshop on Charge-Coupled Devices and Advanced Image Sensors*, pages 157-160, 2005.]

Wich means

from **$5.5 \mu\text{s}$** to **$122.6 \mu\text{s}$** , or from **2200** to **50000 instructions**,
on a 32-bit μp at 400 MHz for any kind of high level spike processing.

Conclusion

- New philosophy of spike processing using a standalone multi-task environment directly connected to the AER bus,
- achieving up to **840 KEvents/s** while constructing a frame and
- letting the execution of [2200, 50000] 32 bit μ p instructions,
- allowing high level spike processing while performing AER communication for high demanding applications.

Acknowledgements

We want to special thank our colleague at Computer Architecture and Technology Department from the University of Seville, Francisco Gómez-Rodríguez, for his great interest and his useful comments when carrying out this work. We would also like to thank the NSF sponsored Telluride Neuromorphic Engineering Workshop, where this idea was born in a discussion group participated by Daniel Fasnacht, Giacomo Indiveri, Alejandro Linares-Barranco and Francisco Gómez-Rodríguez.

This work was supported by Spanish grant TEC2006-11730-C03-02 (SAMANTA 2).



Conclusion

- New philosophy of spike processing using a standalone multi-task environment directly connected to the AER bus,
- achieving up to **840 KEvents/s** while constructing a frame and
- letting the execution of [2200, 50000] 32 bit μ p instructions,
- allowing high level spike processing while performing AER communication for high demanding applications.

Acknowledgements

We want to special thank our colleague at Computer Architecture and Technology Department from the University of Seville, Francisco Gómez-Rodríguez, for his great interest and his useful comments when carrying out this work. We would also like to thank the NSF sponsored Telluride Neuromorphic Engineering Workshop, where this idea was born in a discussion group participated by Daniel Fasnacht, Giacomo Indiveri, Alejandro Linares-Barranco and Francisco Gómez-Rodríguez.

This work was supported by Spanish grant TEC2006-11730-C03-02 (SAMANTA 2).



Conclusion

- New philosophy of spike processing using a standalone multi-task environment directly connected to the AER bus,
- achieving up to **840 KEvents/s** while constructing a frame and
- letting the execution of [**2200, 50000**] 32 bit μ p instructions,
- allowing high level spike processing while performing AER communication for high demanding applications.

Acknowledgements

We want to special thank our colleague at Computer Architecture and Technology Department from the University of Seville, Francisco Gómez-Rodríguez, for his great interest and his useful comments when carrying out this work. We would also like to thank the NSF sponsored Telluride Neuromorphic Engineering Workshop, where this idea was born in a discussion group participated by Daniel Fasnacht, Giacomo Indiveri, Alejandro Linares-Barranco and Francisco Gómez-Rodríguez.

This work was supported by Spanish grant TEC2006-11730-C03-02 (SAMANTA 2).



Conclusion

- New philosophy of spike processing using a standalone multi-task environment directly connected to the AER bus,
- achieving up to **840 KEvents/s** while constructing a frame and
- letting the execution of [**2200, 50000**] 32 bit μ p instructions,
- allowing high level spike processing while performing AER communication for high demanding applications.

Acknowledgements

We want to special thank our colleague at Computer Architecture and Technology Department from the University of Seville, Francisco Gómez-Rodríguez, for his great interest and his useful comments when carrying out this work. We would also like to thank the NSF sponsored Telluride Neuromorphic Engineering Workshop, where this idea was born in a discussion group participated by Daniel Fasnacht, Giacomo Indiveri, Alejandro Linares-Barranco and Francisco Gómez-Rodríguez.

This work was supported by Spanish grant TEC2006-11730-C03-02 (SAMANTA 2).

