

“Exploration of Low Power Adders for a SIMD Data Path”

G. Paci, *University of Bologna and IMEC*


P. Marchal, *IMEC*

L. Benini, *University of Bologna*



University of Bologna

Viale Risorgimento, 2
40136 Bologna
Italy



IMEC

Kapeldreef 75
B-3001 Leuven
Belgium

Market trends

The ambient intelligence market is oriented to produce devices with as many integrated features as possible:

- Media Player (mp3, avi, jpg, mpg...)
- 3D Game
- Web browser
- Video and voice calling
- Wi Fi, GSM, UMTS

These devices are often built to be handheld, therefore the designs constraints are:

- High Performance
- Real time
- Low Power

Time to market and NRE cost

Time to Market and Non-Recurring Engineering (NRE) cost have to be small.

The flexibility and programmability properties allow to meet both constraints.

Unfortunately these features come with a significant energy penalty.

This penalty can be partially recovered by the parallelism and especially by the SIMD (Single Instruction Multiple Data) structure.

Why SIMD is a key

1. Parallelism increases performance
 - It can be exchanged for energy saving (e.g. by voltage scaling)
2. SIMD reduces the cost associated with the instruction decoding

This technique leads high performance with flexibility and programmability but requires software able to exploit them

Hence many related works exist on the benefits *of* and compilation techniques *for* subword parallelism

But the implementation of subword parallel data paths for energy has received limited attention.

SIMD adders analysis

This work is focused on the energy efficiency of various SIMD adders and it analyzes the impact that different templates have with respect to area and performance.

Adders are synthesized with Synopsys Physical compiler using the UMC High Performance 0.13 μ m CMOS standard cell library.

Physical compiler provides:

- Area estimation
- Performance estimation

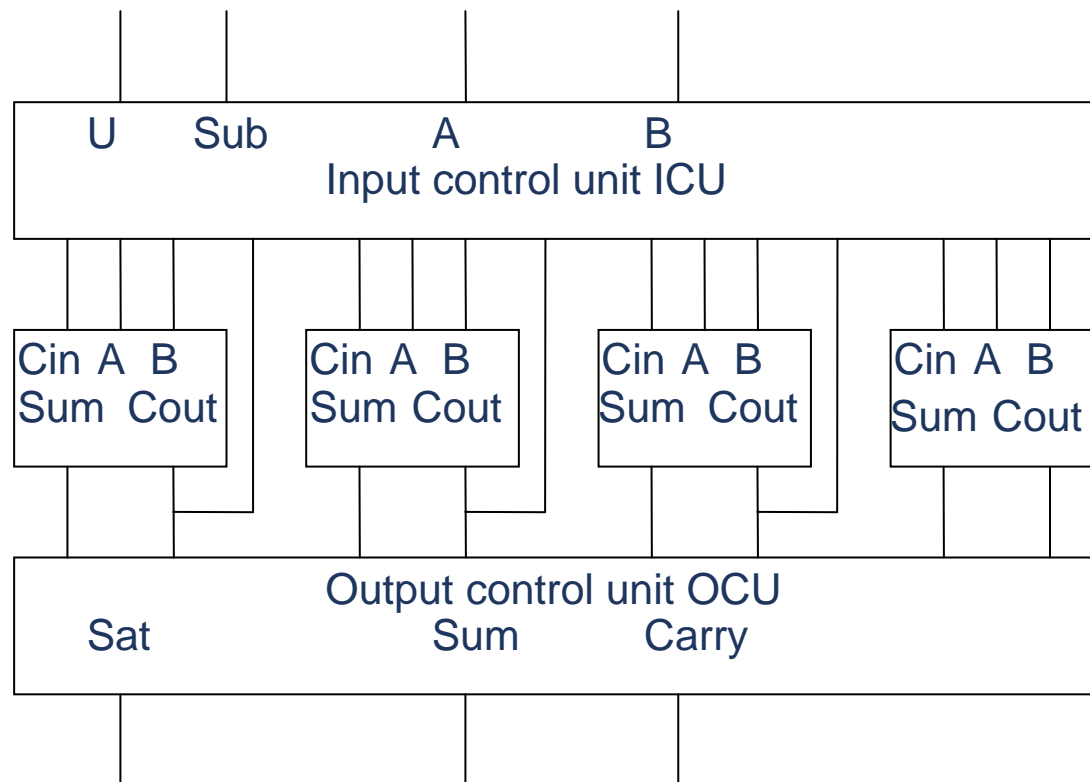
Power compiler provides:

- Power estimation

Typically SIMD adder design

Usually a SIMD adder is designed cascading subunits for the smallest operand length into a larger structure.

Our template for a typical SIMD adder unit is:



Typically SIMD adder design

Synthesis of the previous template sizing the cells to meet the timing constraint for 32bit operation.


Since the other operations (8 bit and 16 bit) have a shorter critical path, they are executed using oversized cell leads a loss in energy efficiency.

Therefore we have compared the power consumption of four 8bit and one 32bit adders both optimized for power.

The analysis compares adders in the Synopsys DesignWare library:

- Ripple carry (rpl)
- Ripple carry select (rpcs)
- Carry look ahead (cla)
- Cary look ahead select (clsa)
- Brent Kung (bk)
- Conditional Sum (csm)

Power related to the operand length and operating frequency



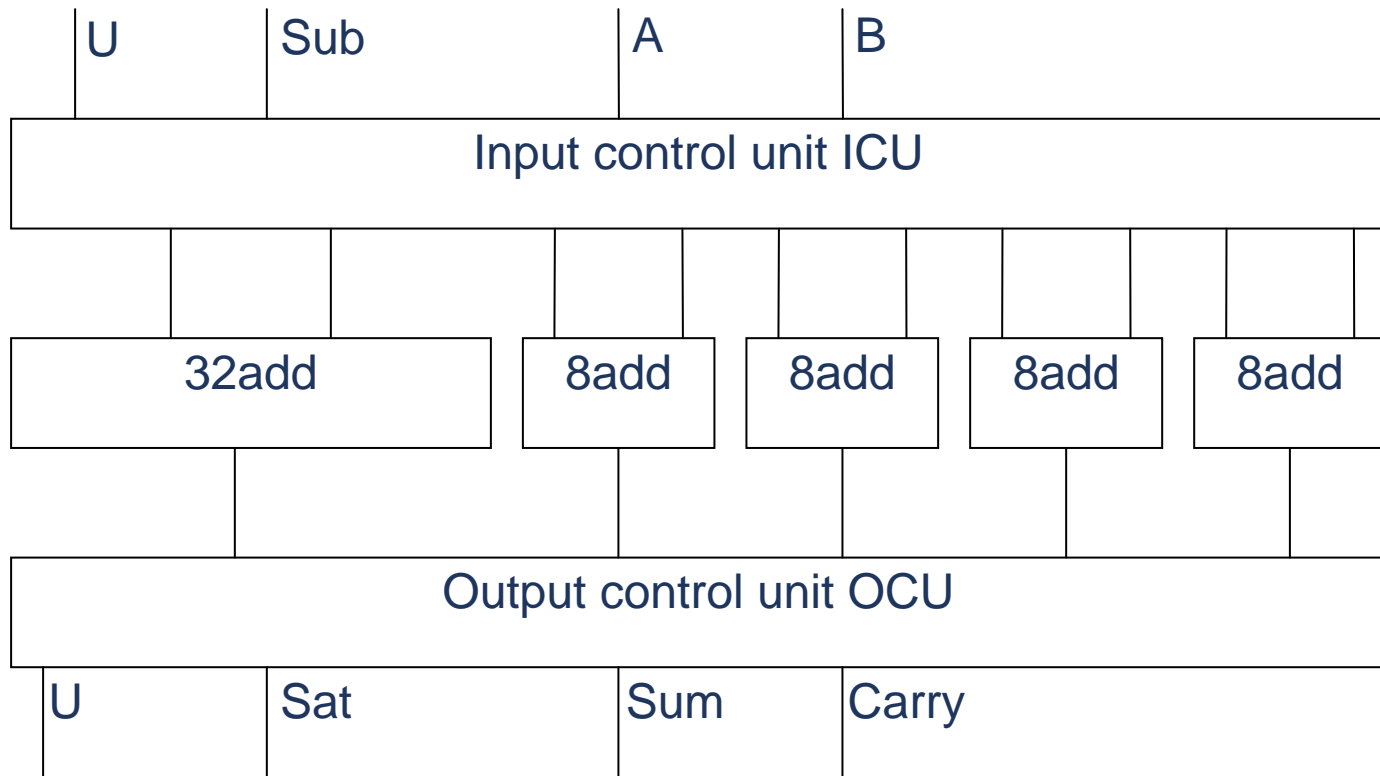
Considering the best solution for specific frequency and data length:

At 100Mhz four 8bit ripple carry adders consume 94 μ W instead of 160 μ W consumed by one 32bit ripple carry select adder. **42% less power**

At 700Mhz four 8bit Brent-Kung adders consume 810 μ W instead of 1221 μ W consumed by one 32bit Brent Kung adder. **34% less power**

Combined SIMD adder

The following template is our SIMD adder that exploits the energy benefit provided by using dedicated adders for each operand length



Combined SIMD adder

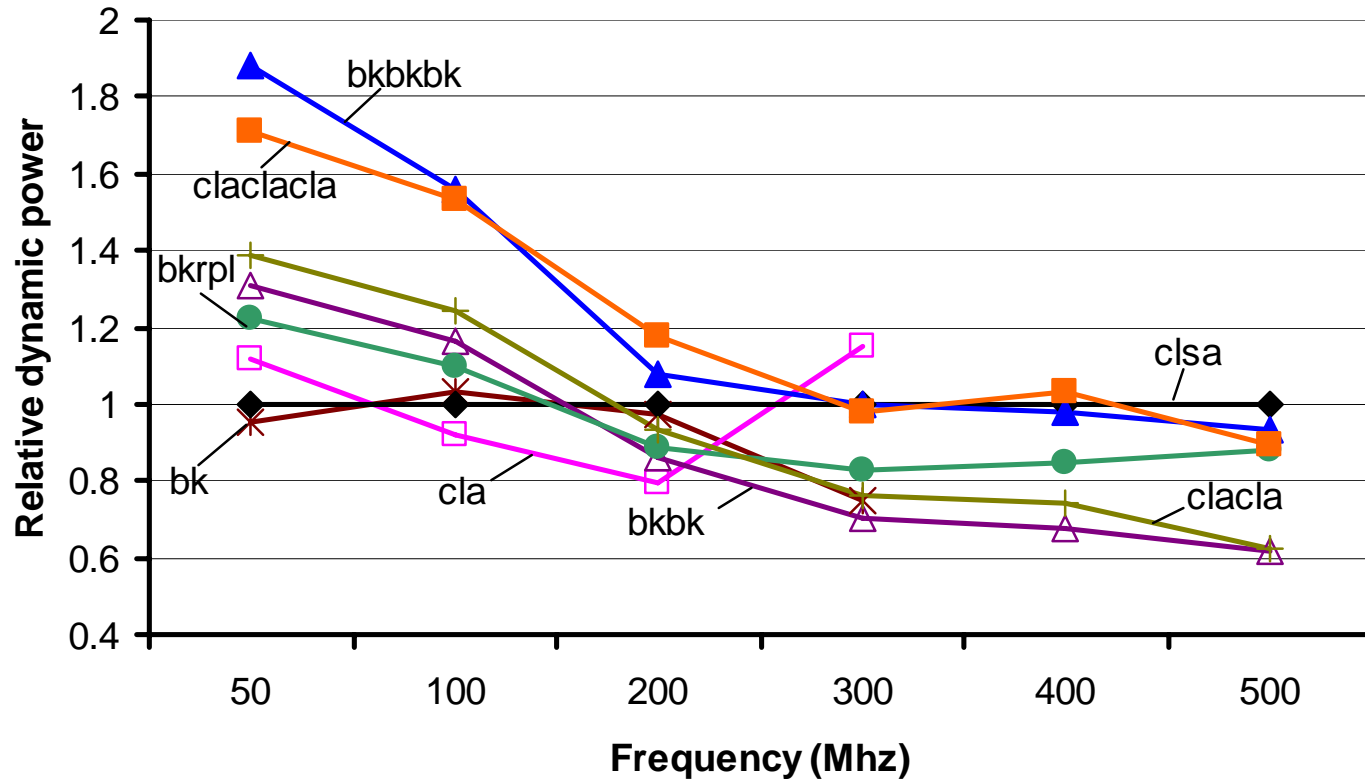
The control logic and the operand isolation being in the combined SIMD introduce an extra power cost.

To analyze the overhead, we have synthesized several combined SIMD adders with different parallelism pattern (4x8bit-1x32bit and 4x8bit-2x16bit-1x32bit) and adder type combination (e.g Brent Kung for 32 bit adder, ripple carry for 8 bit adders).

The typically cascade SIMD adder has been synthesized for various adder types as well.

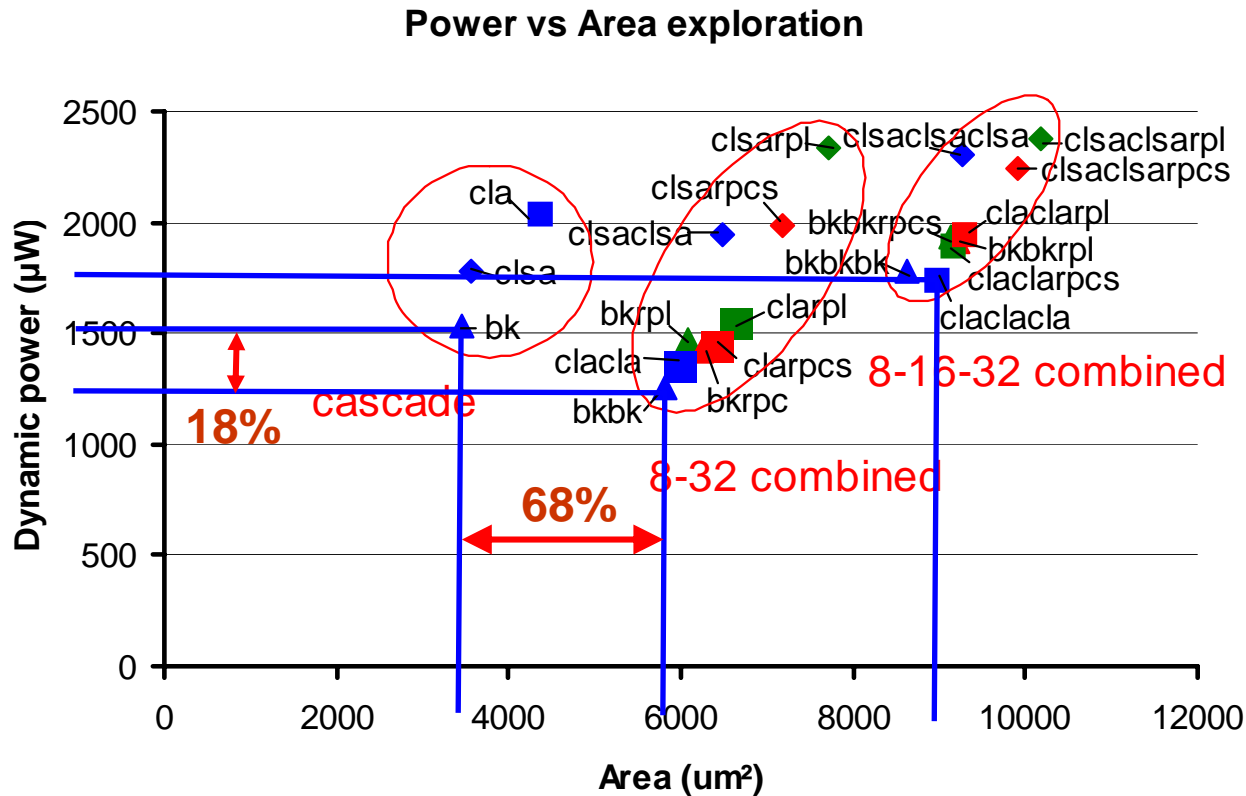
Results

Power vs Frequency relative exploration



Power consumption of combined SIMD adders 32-8bit, combined SIMD adders 32-16-8bit and cascade SIMD normalized to the carry look ahead select cascade SIMD adder. Scenario 50%/50% 8/32bit

Power vs Area exploration



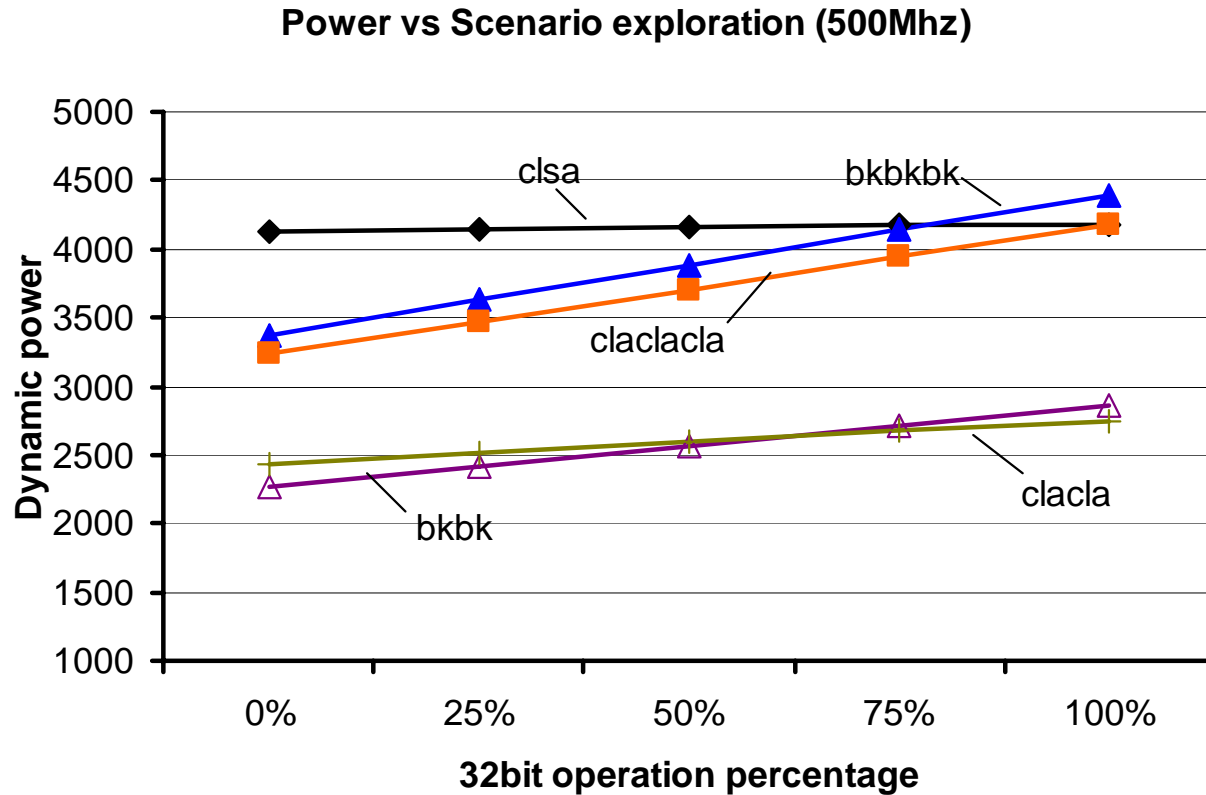
Power and area trade-off exists between the presented SIMD adders. Operating frequency 300Mhz and 50%/50% occurrence of 8/32-bit operand length

Power vs Area exploration

	cascade bk		combined bkbk		combined clacla	
	area	power	area	Power	area	power
adders	2028	914	3724 (1.84)	697 (0.76)	5895 (2.91)	733 (0.80)
ICU	508	207	1327 (2.61)	259 (1.23)	2071 (4.07)	634 (3.06)
OCU	926	410	755 (0.82)	296 (0.72)	1009 (1.09)	376 (0.92)
Total	3462	1531	5806 (1.68)	1252 (0.82)	8975 (2.59)	1743 (1.14)

Area is in μm^2 and power is in μW . Scenario 50%/50% 8-32bit. Operating frequency 300Mhz. The values in the brackets are the relative values compared to the data of the cascaded adder

Power vs Scenario exploration



Power consumption of SIMD adders at 500Mhz is presented for multiple usage scenarios

Conclusion

- The common practice of partitioning a normal adder in sub-units is not the most energy efficient one
- A Better solution is combining adders optimized for each operand length into a single SIMD unit
 - The operand isolation inside the combined SIMD gives a considerable penalty which is recovered only for relative high frequency
 - The area extra cost has to be paid
- The scenario influences the best SIMD solution

Future work

- To extend this work for other operand length (24,64,128 bit)
- Apply the same energy optimization concept to an entire data path

“Thank you”

G. Paci, ***University of Bologna and IMEC***


P. Marchal, ***IMEC***

L. Benini, ***University of Bologna***



University of Bologna

Viale Risorgimento, 2
40136 Bologna
Italy



IMEC

Kapeldreef 75
B-3001 Leuven
Belgium