

# Lucas Kanade Optical Flow Estimation on the TI C66x Digital Signal Processor



Fan Zhang, Yang Gao and Jason D. Bakos



**2014 IEEE High Performance Extreme  
Computing Conference(HPEC '14)**

---

# What is Optical Flow

---

- **Evaluates the pixel movement in video stream**
  - Represented as a dense 2D fields
- **Many applications need to apply real-time optical flow**
  - Robotic vision
  - Augmented reality
- **Computationally intensive**



# TI C66x Multicore DSP

- **Unique architectural features**
  - Eight cores
  - Statically scheduled VLIW/SIMD instructions
  - 2 register files and 8 functional units per core

	<b>Tesla K20X GPU</b>	<b>Intel i7 Ivy Bridge</b>	<b>TI C6678 Keystone</b>	<b>NVIDIA Tegra K1</b>
	Server GPU	Server CPU	<b>Embedded</b>	Embedded
<b>Peak single precision throughput</b>	3.95 Tflops	448 Gflops	<b>160 Gflops</b>	327 Gflops
<b>Peak Power</b>	225 W	77 W	<b>&lt;10 W</b>	<20 W
<b>DRAM bandwidth</b>	250 GB/s	25.6 GB/s	<b>12.8 GB/s</b>	17.1 GB/s

---

# Why using DSP?

---

- Low power consumption
- High performance floating point arithmetic
- Strong potential for computer vision tasks
  - 1D vs 2D (signal processing vs image processing)

# Gradient-based Optical Flow Solver

- Optical flow evaluation

$$f(x, y, t) = f(x + \Delta x, y + \Delta y, t + \Delta t)$$

- First-order approximation

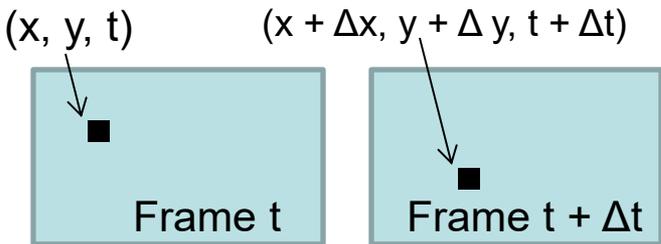
$$f(x, y, t) = f(x, y, t) + \frac{\partial f}{\partial x} \Delta x + \frac{\partial f}{\partial y} \Delta y + \frac{\partial f}{\partial t} \Delta t$$

Gradient of pixel in x, y  
and t direction, known

$$\frac{\partial f}{\partial x} v_x + \frac{\partial f}{\partial y} v_y = -\frac{\partial f}{\partial t}$$

Optical flow, unknown

One  
formula  
with two  
unknowns



# Lucas Kanade Method

If we assume the pixel adjacent to the center has the same optical flow as the center

<b>x-1,y-1</b>	<b>x,y-1</b>	<b>x+1,y-1</b>
<b>x-1,y</b>	<b>x,y</b>	<b>x+1,y</b>
<b>x-1,y+1</b>	<b>x,y+1</b>	<b>x+1,y+1</b>

$$\begin{cases} \frac{\partial f(x-1, y-1)}{\partial x} v_x + \frac{\partial f(x-1, y-1)}{\partial y} v_y = -\frac{\partial f(x-1, y-1)}{\partial t} \\ \frac{\partial f(x+1, y+1)}{\partial x} v_x + \frac{\partial f(x+1, y+1)}{\partial y} v_y = -\frac{\partial f(x+1, y+1)}{\partial t} \end{cases}$$

Let  $A = \begin{bmatrix} \frac{\partial f(x-1, y-1)}{\partial x} & \frac{\partial f(x-1, y-1)}{\partial y} \\ \frac{\partial f(x+1, y+1)}{\partial x} & \frac{\partial f(x+1, y+1)}{\partial y} \end{bmatrix}$   $b = \begin{bmatrix} \frac{\partial f(x-1, y-1)}{\partial t} \\ \frac{\partial f(x+1, y+1)}{\partial t} \end{bmatrix}$

$$\begin{bmatrix} V_x \\ V_y \end{bmatrix} = (A^T A)^{-1} A^T b = \begin{bmatrix} \sum \frac{\partial f^2}{\partial x} & \sum \frac{\partial f}{\partial x} \frac{\partial f}{\partial y} \\ \sum \frac{\partial f}{\partial x} \frac{\partial f}{\partial y} & \sum \frac{\partial f^2}{\partial y} \end{bmatrix} \begin{bmatrix} -\sum \frac{\partial f}{\partial x} \frac{\partial f}{\partial t} \\ -\sum \frac{\partial f}{\partial y} \frac{\partial f}{\partial t} \end{bmatrix}$$

Least Square Method



# Image Derivative Computation

$$\frac{\partial f}{\partial x} = \begin{array}{c} A \quad \begin{array}{|c|c|} \hline \text{yellow} & \text{light blue} \\ \hline \end{array} \quad B \\ C \quad \begin{array}{|c|c|} \hline \text{light blue} & \text{light blue} \\ \hline \end{array} \quad D \\ \text{frame } n \end{array} \quad (A - B + C - D) / 2$$

$$\frac{\partial f}{\partial y} = \begin{array}{c} A \quad \begin{array}{|c|c|} \hline \text{yellow} & \text{light blue} \\ \hline \end{array} \quad B \\ C \quad \begin{array}{|c|c|} \hline \text{light blue} & \text{light blue} \\ \hline \end{array} \quad D \\ \text{frame } n \end{array} \quad (A - C + B - D) / 2$$

$$\frac{\partial f}{\partial t} = \begin{array}{c} A \quad \begin{array}{|c|} \hline \text{yellow} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \text{light blue} \\ \hline \end{array} \quad B \\ \text{frame } n \quad \text{frame } n+1 \end{array} \quad A - B$$



---

# Lucas Kanade Method

---

- Input: Frame  $n$ , frame  $n+1$ , window size  $w$
- Output: Optical flow field  $F$
  
- For each pixel  $x, y$ 
  - Compute  $x, y, t$  derivative matrices  $D_x, D_y, D_t$  for its neighbor window
  - Compute optical flow by the least square method

# Derivative Computation (Example)

Image n

10	10	10	10
10	30	10	10
10	10	10	10
10	10	10	10

Image n + 1

10	10	10	10
10	10	10	10
10	10	30	10
10	10	10	10

-10	10	0
-10	10	0
0	0	0

$$\frac{\partial f}{\partial x}$$

-10	-10	0
10	10	0
0	0	0

$$\frac{\partial f}{\partial y}$$

0	0	0
0	20	0
0	0	-20

$$\frac{\partial f}{\partial t}$$



# Least Square Method Example (Neighbor Window Size = 3)

$$\frac{\partial f}{\partial x} \begin{array}{|c|c|c|} \hline -10 & 10 & 0 \\ \hline -10 & 10 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array}$$

$$\frac{\partial f}{\partial y} \begin{array}{|c|c|c|} \hline -10 & -10 & 0 \\ \hline 10 & 10 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array}$$

$$\frac{\partial f}{\partial t} \begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 20 & 0 \\ \hline 0 & 0 & -20 \\ \hline \end{array}$$

$$\sum_w \left( \frac{\partial f}{\partial x} \right)^2 = 400$$

$$\sum_w \left( \frac{\partial f}{\partial y} \frac{\partial f}{\partial t} \right) = 200$$

$$\sum_w \left( \frac{\partial f}{\partial x} \frac{\partial f}{\partial y} \right) = 0$$

$$\sum_w \left( \frac{\partial f}{\partial y} \right)^2 = 400$$

$$\sum_w \left( \frac{\partial f}{\partial x} \frac{\partial f}{\partial t} \right) = 200$$

$$\begin{bmatrix} V_x \\ V_y \end{bmatrix} = \begin{bmatrix} 400 & 0 \\ 0 & 400 \end{bmatrix}^{-1} \begin{bmatrix} -200 \\ -200 \end{bmatrix}$$



---

# Optimize Lucas Kanade Method on DSP

---

- **Derivative computation**
  - SIMD arithmetic
  - Data interleaving
- **Least square method**
  - Loop unrolling
  - SIMD load & arithmetic

# Derivative Computation

10	10	10	10
10	30	10	10
10	10	10	10
10	10	10	10

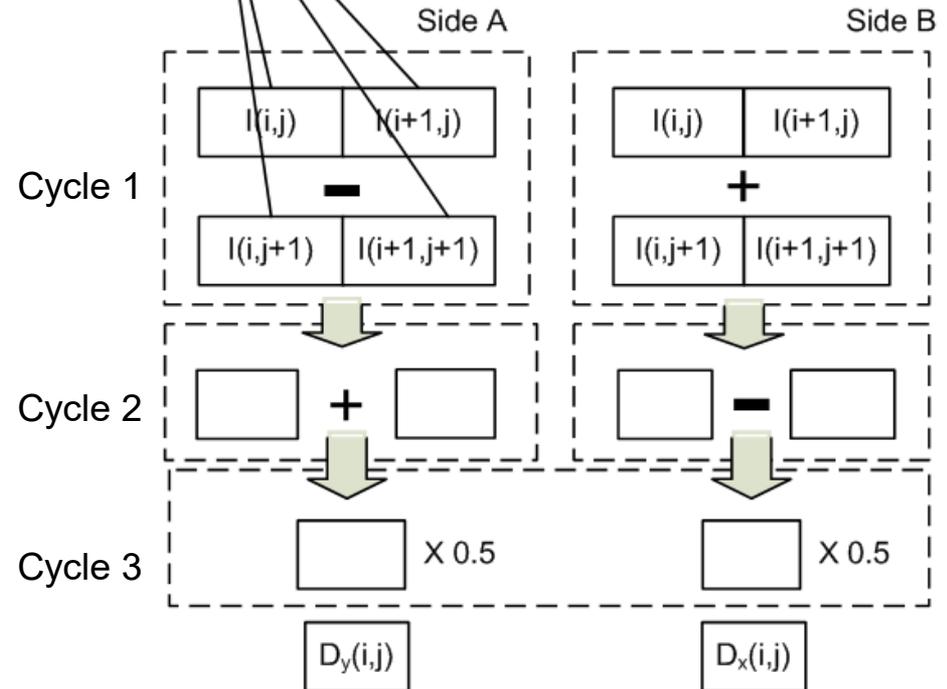
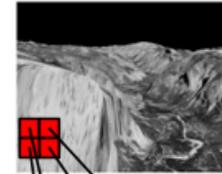
Derivative Computation (Dx, Dy)

-10	10	0
-10	10	0
0	0	0

-10	-10	0
10	10	0
0	0	0

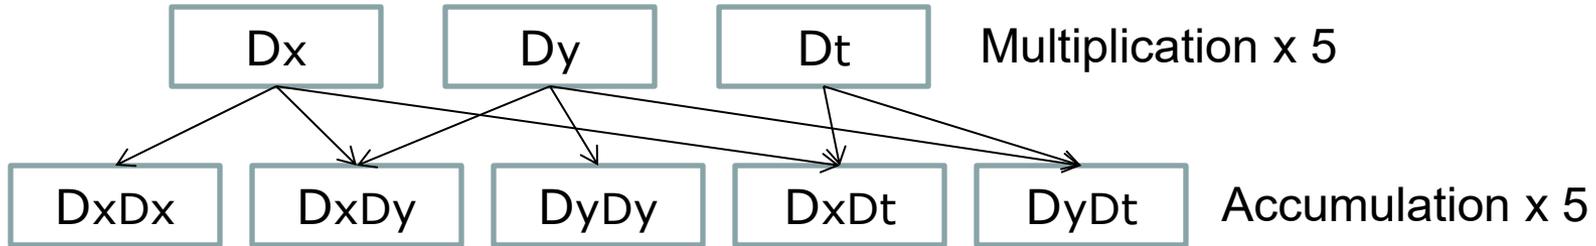
Interleave

-10	-10	10	-10	0	0
-10	10	10	10	0	0
0	0	0	0	0	0

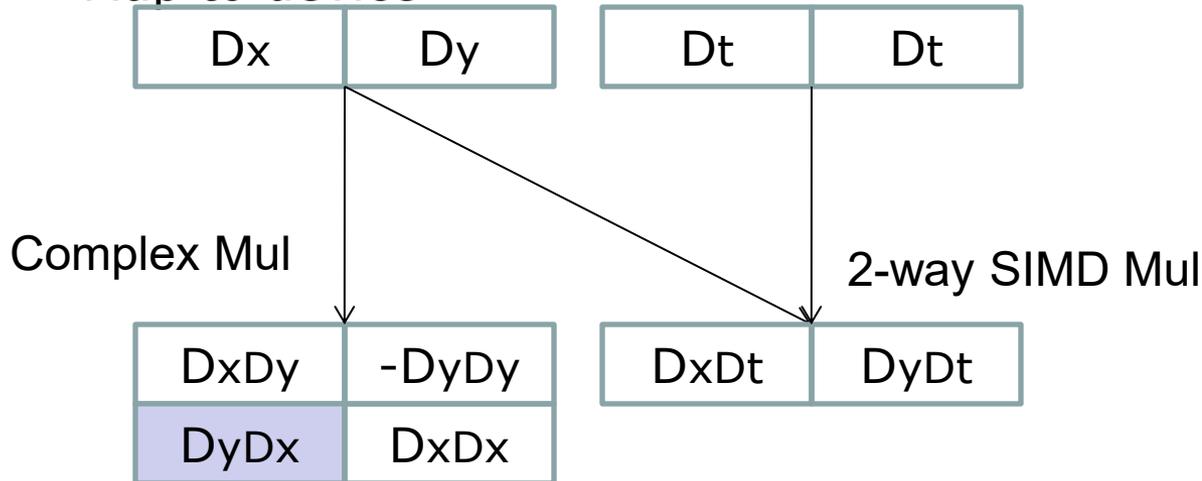


# Least Square Method

- Required computations



- Map to device



$$(a+bj)(c+dj) = (ac-bd) + (ad+bc)j$$



# Least Square Method

- Unrolled 2x
  - Group up loads into SIMD operation

Load Dx Dy

Load Dt

Complex MUL

SIMD MUL

SIMD ADD

DxDy

Dt

-10	-10	10	-10	0	0
-10	10	10	10	0	0
0	0	0	0	0	0

0	0	0
0	20	0
0	0	-20

(Dx, Dy, Dt)

-10	-10	0
-----	-----	---

(DxDx, Dx Dy, Dy Dy, Dx Dt, Dy Dt)

100	100	100	0	0
-----	-----	-----	---	---

(Dx, Dy, Dt)

10	-10	0
----	-----	---

(DxDx, Dx Dy, Dy Dy, Dx Dt, Dy Dt)

100	-100	100	0	0
-----	------	-----	---	---

+

200
0
200
0
0

# Loop Flattening

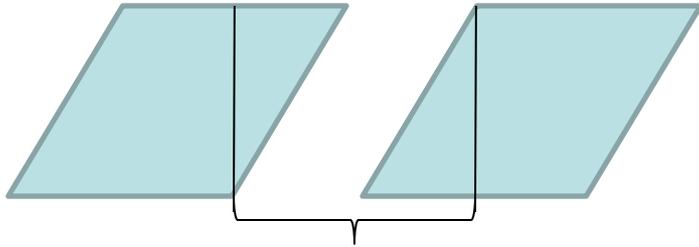
- Software Pipelining

- TI's compiler technique to allow independent loop iterations be executed in parallel
- The consecutive iterations are packed and executed together so that the arithmetic functional units are fully utilized

```
for (i = 0; i < m; ++i)  
  for (j = 0; j < n; ++j)
```

j = 0

j = 1



Pipeline prologue/epilogue overhead

```
for (k = 0; k < m * n; ++k)
```

...

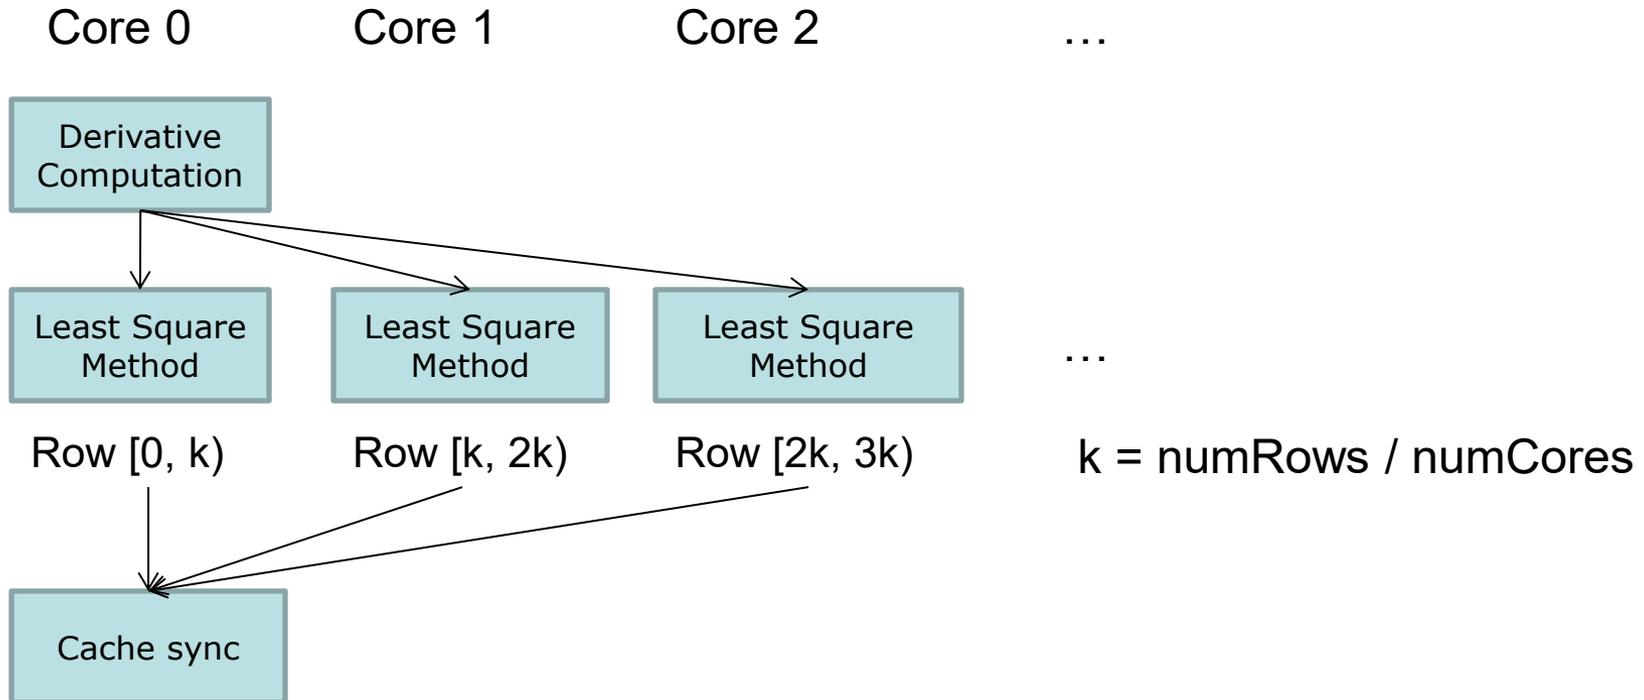
Update i, j;

j = 0

j = 1



# Multicore Utilization



---

# Accuracy Improvement

---

- **Cannot catch movement larger than window size**
  - Gaussian Pyramid
    - A coarse to fine strategy for optical flow computation
    - Catches large movements
- **First order approximation may not be accurate**
  - Iterative refinement
    - Iteratively process and offset pixels until the computed optical flow is converged
    - Introduce data random access



---

# Experiment Setup

---

<b>Platform</b>	<b>#Cores</b>	<b>Implementation</b>	<b>Power Measurement</b>
TI C6678 DSP	8	Our Implementation	TI GPIO-USB Module
ARM Cortex A9	2	Our Implementation	YOKOGAWA WT500 Power Analyzer
Intel i7-2600	4	Our Implementation	Intel RAPL
Tesla K20 GPU	2688	OpenCV	NVIDIA SMI

# Results and Analysis

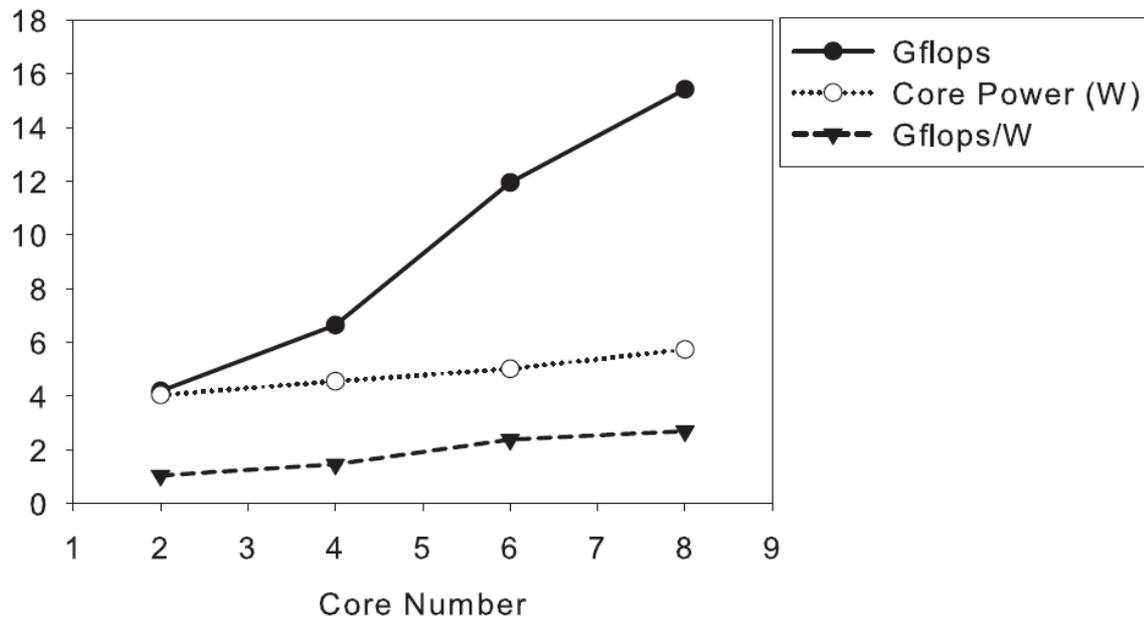
- Highest Performance
  - Tesla K20
- Lowest Power
  - Cortex A9
- Best Power Efficiency
  - TI C66x DSP

Platform	C66x	CortexA9	Intel i7-2600	K20
Actual Gflops/ Peak Gflops	<b>12%</b>	7%	4%	3%
Gflops	15.4	0.7	17.1	<b>108.6</b>
Power (W)	5.7	<b>4.8</b>	52.5	79.0
Gflops/W	<b>2.69</b>	0.2	0.3	1.4

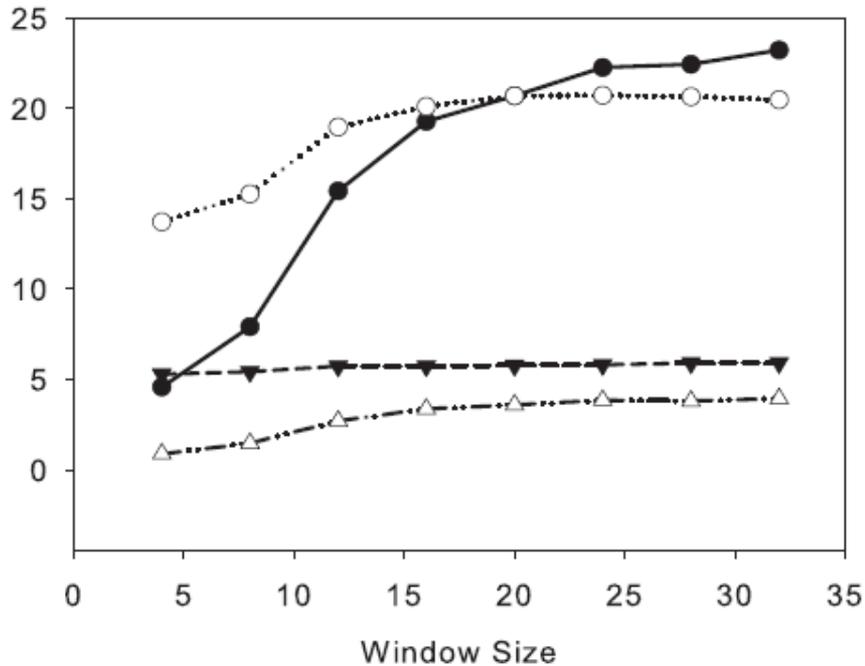


# Results and Analysis

- We achieve linear scalability on multi-cores
- The power efficiency of the DSP is low when its cores are partially utilized
  - Static power consumption



# Results and Analysis



- Performance are related with window size
  - Software pipeline performance
- Loop flattening is able to improve performance significantly on small window size



---

# Conclusion

---

- First research work on DPS accelerated Lucas Kanade method
- Achieve higher energy efficiency and device utilization than GPU and CPU

---

# Q & A

---



# Kernels of Pyramidal Lucas Kanade Method

- Gaussian Blur 28 flop/pixel
  - Derivative Computation 8 flop/pixel
  - Least Square Method  $10^5$  flop/pixel
  - Flow Field Bilinear Interpolation 3 flop/pixel
- Window Size = 16,  
Pyramid Level = 4,  
Iteration = 10
- 