



- Map
- Map + Local reduction
- Reduction
- Scan

- LUT Application
- Local features extraction
- Histogram
- Integral images

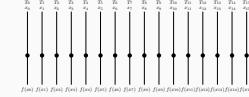
```
void plus_one(int* a, int size, int k) {
    int i = blockDim.x * blockIdx.x +
            threadIdx.x + k;
    if (i < size)
        a[i] = a[i] + 1;
}
```

What do you think about K's impact of the performance ?

1. GPU and architectures (2h, Friday AM)
2. Programming GPUs with CUDA (2h, Friday PM)
3. TP 00 CUDA (Getting started) (3h, Monday or Tuesday)
4. Efficient programming with GPU (part 1) (2h, Wednesday AM)
5. TP 01 CUDA (Mandelbrot) (3h, Friday AM or PM)
6. Efficient programming with GPU (part 2) (2h, Monday 25th)
7. Assignments (remote, async, end of Oct. + Nov. 5th)
8. TP01 Q/A (3h, Friday, Nov. 5th)

Map replicates a function over every element of an index set
The computation of each pixel is independant wrt the others.

out(x,y) = f(in(x,y))



Nothing complicated but take care of memory access pattern.

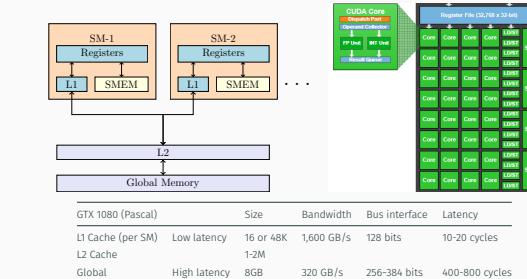
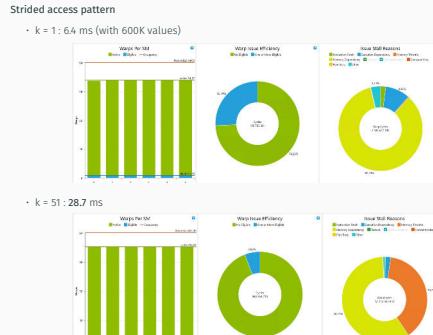
```
void plus_one(int* a, int size, int k) {
    int i = blockDim.x * blockIdx.x +
            threadIdx.x + k;
    if (i < size)
        a[i] = a[i] + 1;
}
```

What do you think about K's impact of the performance ?

Effective Bandwidth vs. Offset for Single Precision

Effective Bandwidth vs. Stride for Single Precision

- Linear sequential access with offset (left) →
- Strided access →

Two types of global memory loads : **Cached** or **Uncached** (L1 disabled)

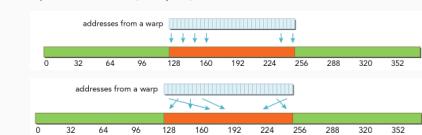
Aligned vs Misaligned

A load is aligned if the first address of a memory access is multiple of 32 bytes

- Memory addresses must be type-aligned (ie `sizeof(T)`)
- Otherwise : poor perf (unaligned load)
- `cudaMalloc` = alignment on 256 bits (at least)

We need a load strategy :

- 32 threads of warp access a 32-bit word = 128 bytes
- 128 bytes = L1 bus width (single load - bus utilization = 100%)
- Access permutation has no (or very low) overhead



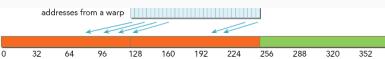
A load is coalesced if a warp accesses a contiguous chunk of data

- Minimize memory accesses caused by warp threads
- Remind : all threads of a warp executes the same instruction
→ if a load, may be 32 different addresses

Misaligned cached loads from L1

- If data are not 128-bits aligned, two loads are required

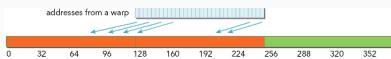
Addresses 96-224 required... but 0-256 loaded



Misaligned cached loads from L1

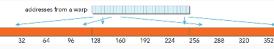
- If data are not 128-bits aligned, two loads are required

Addresses 96-224 required... but 0-256 loaded



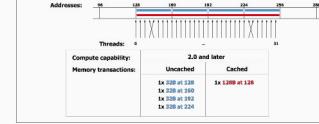
- If data is accessed strided

e.g. $u[2*k], \dots$

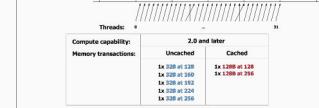


Coalesced Memory Access (Summary 1)

Aligned accesses (sequential/non-sequential)

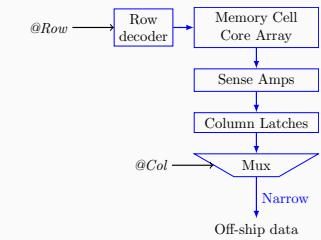


Mis-aligned accesses (sequential/non-sequential)



How memory works for real

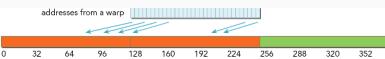
- DRAM is organised in 2D Core array
- Each DRAM core array has about 16M bits



Misaligned cached loads from L1

- If data are not 128-bits aligned, two loads are required

Addresses 96-224 required... but 0-256 loaded



- If data is accessed strided

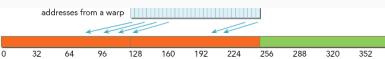
e.g. $u[2*k], \dots$



Loads from global (uncached) memory

- If data are not 128-bits aligned, two loads are required

Addresses 96-224 required... but 0-256 loaded

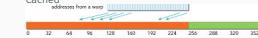
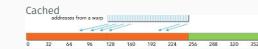


- If data is accessed strided

e.g. $u[2*k], \dots$

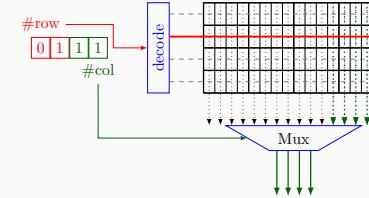


Same idea but memory is split in segments of 32 bytes



Example

- A 4 x 4 memory cell
- With 4 bits pin interface width



DRAM Burst

Reading from a cell in the core array is a very slow process ($1/N$ th of the interface speed):

- DDR : Core speed = $\frac{1}{2}$ interface speed
- DDR2/GDDR3 : Core speed = $\frac{1}{4}$ interface speed
- DDR3/GDDR4 : Core speed = $\frac{1}{8}$ interface speed

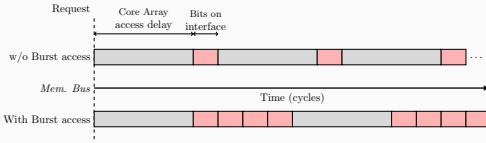
DRAM Burst

Reading from a cell in the core array is a very slow process ($1/N$ th of the interface speed):

- DDR : Core speed = $\frac{1}{2}$ interface speed
- DDR2/GDDR3 : Core speed = $\frac{1}{4}$ interface speed
- DDR3/GDDR4 : Core speed = $\frac{1}{8}$ interface speed

Solution : Bursting

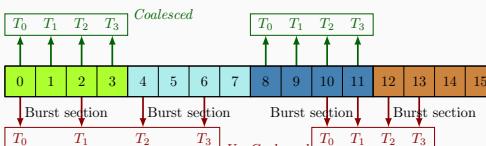
Load N x interface width of the same row



Better, but not enough to saturate the memory bus (will see later a solution).

Summary

- Use coalesced (contiguous and aligned) access to memory :



- If all threads of a warp execute a load instruction into the same burst section \rightarrow only one DRAM request

- Otherwise :
 - Multiple DRAM requests are made
 - Some bytes transferred are not used

18

Using shared memory

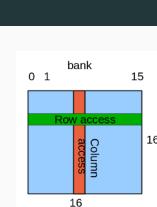
Transposition

```
int x = blockIdx.x * blockDim.x + threadIdx.x;
```

```
int y = blockIdx.y * blockDim.y + threadIdx.y;
```

```
// transpose with boundary test
```

```
if (x < w && y < h)
  out[x * width + y] = in[y * width + x]
```



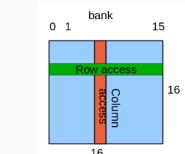
Transposition

```
int x = blockIdx.x * blockDim.x + threadIdx.x;
```

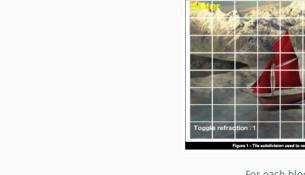
```
int y = blockIdx.y * blockDim.y + threadIdx.y;
```

```
// transpose with boundary test
```

```
if (x < w && y < h)
  out[x * width + y] = in[y * width + x]
```



Tiling and memory privatization in shared memory

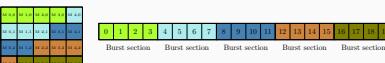


For each block:

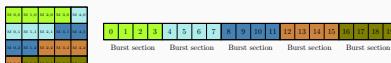
- read the tile from global to private block memory
- process the block
- write the tile from the private block memory to global memory

```
void mykernel() {
  __shared__ float private_mem[TILE_WIDTH][TILE_WIDTH];
}
```

Q : how to make coalesced loads with 2D-arrays?



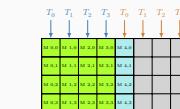
Q : how to make coalesced loads with 2D-arrays?



- Add padding to align rows on 256-bits boundaries

- Thread reads data column-wise

```
ima(tid.x, tid.y) = tid.y * pitch + tid.x
```



Where are non-coalesced access?

$\rightarrow a[x][y]$

- Reads are coalesced

- Write are strided

19

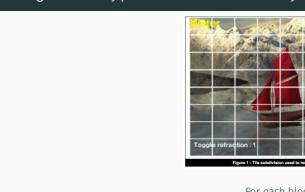
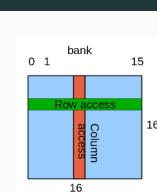
Transposition

```
int x = blockIdx.x * blockDim.x + threadIdx.x;
```

```
int y = blockIdx.y * blockDim.y + threadIdx.y;
```

```
// transpose with boundary test
```

```
if (x < w && y < h)
  out[x * width + y] = in[y * width + x]
```



For each block:

- read the tile from global to private block memory
- process the block
- write the tile from the private block memory to global memory

```
void mykernel() {
  __shared__ float private_mem[TILE_WIDTH][TILE_WIDTH];
}
```

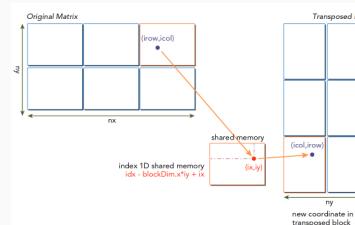
19

Collaborative loading and writing when BLOCKDIM = TILEDIM

- All threads load one or more data
 - Access must be coalesced
 - Use barrier synchronization to make sure that all threads are ready to start the phase
- ```
void tiledKernel(unsigned char * in, unsigned char * out, int w, int h) {
 __shared__ float tile[TILE_WIDTH][TILE_WIDTH];
 int block_ptr = in + ...; // Compute pointer to the beginning of the tile
 for (int y = threadIdx.y + blockDim.y; y < TILE_WIDTH; y += blockDim.y)
 for (int x = threadIdx.x; x < TILE_WIDTH; x += blockDim.x)
 if (x < w && y < h)
 tile[threadIdx.y][threadIdx.x] = in[y * pitch + x];
 __syncthreads();
 __syncthreads();
 __Write
 if (x < w && y < h)
 out[y * pitch + x] = tile[threadIdx.y][threadIdx.x];
}
```

## Collaborative loading when BLOCKDIM < TILEDIM

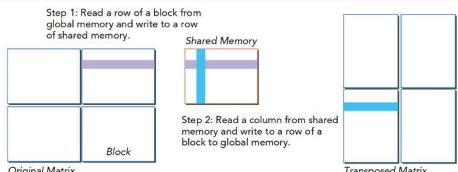
```
{
 __shared__ float tile[TILE_WIDTH][TILE_WIDTH];
 int block_ptr = in + ...; // Compute pointer to the beginning of the tile
 for (int y = threadIdx.y + blockDim.y; y < TILE_WIDTH; y += blockDim.y)
 for (int x = threadIdx.x; x < TILE_WIDTH; x += blockDim.x)
 if (x < width && y < height)
 tile[y][x] = block_ptr[y * pitch + x];
 __syncthreads();
}
```



Performance (GB/s on TESLA K40)

|  | Copy (baseline) | Transpose Naive | Transpose Tiled |
|--|-----------------|-----------------|-----------------|
|  | 17715 GB        | 68.98           | 116.82          |

Tiled transposition in shared memory (1/2)



Tiled transposition in shared memory (2/2)

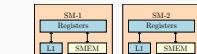
```
void tiledKernel(unsigned char * in, unsigned char * out, int w, int h) {
 __shared__ float tile[TILE_WIDTH][TILE_WIDTH];
 int block_ptr = in + ...; // src
 int Y = threadIdx.y + blockDim.y; // src
 int X = threadIdx.x + blockDim.x; // dst
 int Y = threadIdx.y + blockDim.y + blockDim.x; // dst
 int X = threadIdx.x; // dst

 // Load a line
 if (x < w && y < h)
 tile[threadIdx.y][threadIdx.x] = in[y * pitch + x];

 __syncthreads();
 // Write a line from a column in private mem
 if (x < w && y < h)
 out[Y * pitch + X] = tile[threadIdx.y][threadIdx.x];
}
```

24

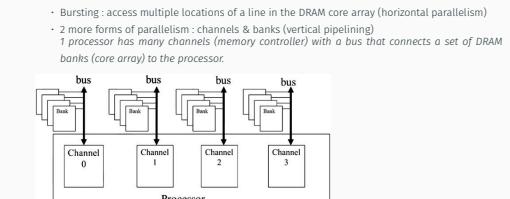
About shared memory



|                   | GTx 1080 (Pascal) | Size      | Bandwidth  | Latency        |
|-------------------|-------------------|-----------|------------|----------------|
| L1 Cache (per SM) | Low latency       | 16 or 48K | 1,600 GB/s | 10-20 cycles   |
| L2 Cache          |                   | 1-2M      |            |                |
| Global            | High latency      | 8GB       | 320 GB/s   | 400-800 cycles |

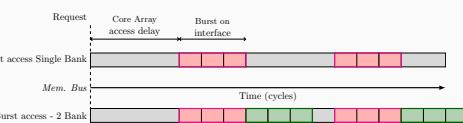
25

DRAM Banks



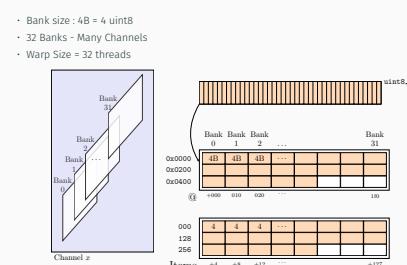
28

Bank conflicts in shared memory

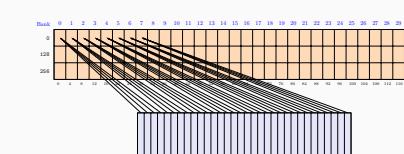


30

Concrete Example for Shared Memory

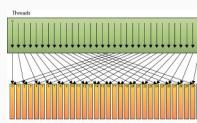


31

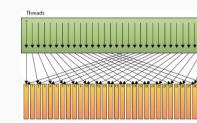


33

### 2-way conflicts



### 2-way conflicts

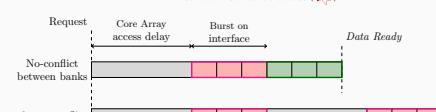


Conflict = Serialized access (red)

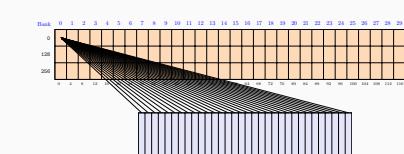
32



32

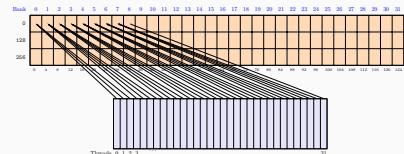


34

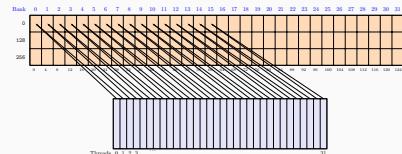


| op                | Items        | Bank used   | Conflict Free | #Loads |
|-------------------|--------------|-------------|---------------|--------|
| load M[tid.x]     | [0,1,...,31] | [0,1,...,7] | ✓             | 1x8    |
| load M[tid.x % 4] |              |             |               |        |

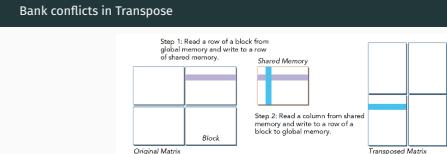
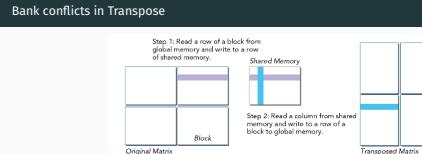
34



| op                | Items                         | Bank used      | Conflict Free | #Loads |
|-------------------|-------------------------------|----------------|---------------|--------|
| load M[tid.x]     | [0, 1, ..., 31]               | [0, 1, ..., 7] | ✓             | 1x8    |
| load M[tid.x % 4] | [0, 1, 2, 3, 0, 1, 2, 3, ...] | [0]            |               | 1x1    |
| load M[tid.x + 1] | [1, 2, 3, ..., 32]            | [0, 1, ..., 8] | ✓             | 1x9    |
| load M[tid.x * 2] |                               |                |               |        |



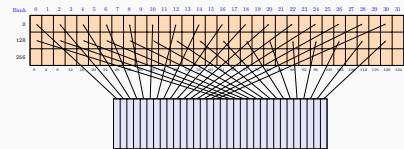
| op                 | Items                         | Bank used       | Conflict Free | #Loads |
|--------------------|-------------------------------|-----------------|---------------|--------|
| load M[tid.x]      | [0, 1, ..., 31]               | [0, 1, ..., 7]  | ✓             | 1x8    |
| load M[tid.x % 4]  | [0, 1, 2, 3, 0, 1, 2, 3, ...] | [0]             |               | 1x1    |
| load M[tid.x + 1]  | [1, 2, 3, ..., 32]            | [0, 1, ..., 8]  | ✓             | 1x9    |
| load M[tid.x * 2]  | [0, 2, 4, ..., 62]            | [0, 1, ..., 15] | ✓             | 1x16   |
| load M[tid.x * 8]  | [0, 8, ..., 248]              | [0, 2, ..., 30] | X             | 2x16   |
| load M[tid.x * 12] |                               |                 |               |        |



```
_shared__ a[16][16];
Y = by*16+ty;
X = bx*16+tx;
```

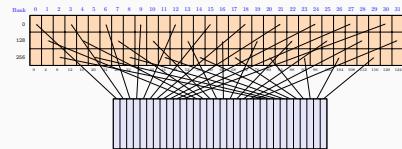
```
a[y][x] = A[Y][X]
__syncthreads()
B[Y][X] = a[x][y];
```

Reading a column may create bank conflicts



| op                 | Items                         | Bank used       | Conflict Free | #Loads |
|--------------------|-------------------------------|-----------------|---------------|--------|
| load M[tid.x]      | [0, 1, ..., 31]               | [0, 1, ..., 7]  | ✓             | 1x8    |
| load M[tid.x % 4]  | [0, 1, 2, 3, 0, 1, 2, 3, ...] | [0]             |               | 1x1    |
| load M[tid.x + 1]  | [1, 2, 3, ..., 32]            | [0, 1, ..., 8]  | ✓             | 1x9    |
| load M[tid.x * 2]  | [0, 2, 4, ..., 62]            | [0, 1, ..., 15] | ✓             | 1x16   |
| load M[tid.x * 8]  | [0, 8, ..., 248]              | [0, 2, ..., 30] | X             | 2x16   |
| load M[tid.x * 12] |                               |                 |               |        |

34



| op                 | Items                         | Bank used       | Conflict Free | #Loads |
|--------------------|-------------------------------|-----------------|---------------|--------|
| load M[tid.x]      | [0, 1, ..., 31]               | [0, 1, ..., 7]  | ✓             | 1x8    |
| load M[tid.x % 4]  | [0, 1, 2, 3, 0, 1, 2, 3, ...] | [0]             |               | 1x1    |
| load M[tid.x + 1]  | [1, 2, 3, ..., 32]            | [0, 1, ..., 8]  | ✓             | 1x9    |
| load M[tid.x * 2]  | [0, 2, 4, ..., 62]            | [0, 1, ..., 15] | ✓             | 1x16   |
| load M[tid.x * 8]  | [0, 8, ..., 248]              | [0, 2, ..., 30] | X             | 2x16   |
| load M[tid.x * 12] | [0, 12, ..., 372]             | [0, 1, ..., 31] | ✓             | 1x32   |

34

### Solution to bank conflicts

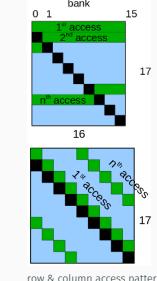
- With padding (to WRAP\_SIZE + 1)

```
_shared__ a[17][17];
Y = by*16+ty;
X = bx*16+tx;
```

```
a[y][x] = A[Y][X]
```

```
__syncthreads()
```

```
B[Y][X] = a[x][y];
```



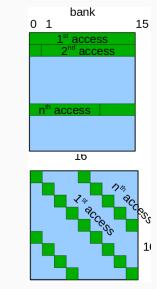
row & column access pattern

```
_shared__ a[16][16];
Y = by*16+ty;
X = bx*16+tx;
```

```
a[f(x,y)] = A[Y][X]
```

```
__syncthreads()
```

```
B[Y][X] = a[f(y,x)]
```



row & column access pattern

### Naive Stencil Implementation

Local average with a rectangle of radius  $r$ . (Ignoring border problems for now).

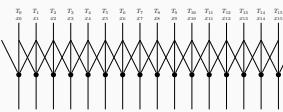
```
_global void boxfilter(const int* in, int* out, int w, int h, int r)
{
 int x = blockIdx.x * blockDim.x + threadIdx.x;
 int y = blockIdx.y * blockDim.y + threadIdx.y;
 if (x < r || x >= w - r) return;
 if (y < r || y >= h - r) return;

 int sum = 0;
 for (int kx = -r; kx <= r; ++kx)
 for (int ky = -r; ky <= r; ++ky)
 sum += in[(y+ky)*w + (x+kx)];
 out[y*w + x] = sum / ((2*r+1) * (2*r+1));
}
```

The computation of a single pixel relies on its neighbors

Use case :

- Dilation/Erosion
- Box (Mean) / Convolution Filters
- Bilateral Filter
- Gaussian Filter
- Sobel Filter



### Performance (GB/s on TESLA K40)

| Copy (baseline) | Transpose Naive | Transpose Tiled | Transpose Tiled+Pad |
|-----------------|-----------------|-----------------|---------------------|
| 17715 GB        | 68.98           | 116.82          | 121.83              |

38

- Super fast access (almost as fast as registers)
- But limited ressources (64-96Kb by SM)

Use it carefully to avoid reducing the occupancy.

39

### Naive Stencil Performance

- Let's say, we have this GPU :

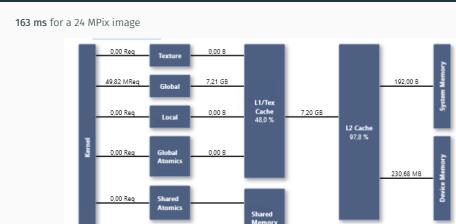
Peak power : 1500 GFlops and Memory Bandwidth : 200 GB/s

- All threads access global memory

- 1 Memory access for 1 FP Addition
- Requires  $1500 \times \text{size}(f) = 6 \text{ TB/s}$  of data
- But only 200 GB/s mem bandwidth  $\rightarrow 50 \text{ GFLOPS}$  (3% of the peak)

### Compute-to-global-memory-access ratio

We need to have  $\#FLOP / \#GlobalMemAccess \geq 30$  to reach the peak



• Problem : too many access to global memory

Shared Memory ↗ → Number of ACTIVE Warp / SM ↘

### Occupancy

number of warps executed at the same time divided by the maximum number of warps that can be executed at the same time.

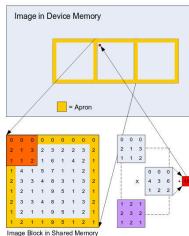
| Generation    | Warps per SM | Warps per scheduler | Active threads limits |
|---------------|--------------|---------------------|-----------------------|
| Maxwell (5.2) | 64           | 16                  | 2048                  |
| Pascal (6.1)  | 64           | 16                  | 2048                  |
| Volta (7.0)   | 64           | 16                  | 1024                  |
| Turing (7.5)  | 32           | 8                   | 1024                  |

40

### Non-local Access Pattern with Tiling

## Handling Border

- Solution : tiling, copy data to shared memory per block first.

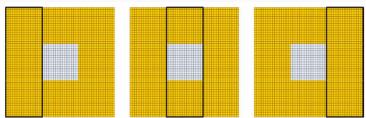


1. Add border to the image to have in-memory access
2. Copy tile + border to shared memory

45

46

1. The bad way : each thread copies one value and border threads are then idle.



2. The good way : a thread may copy several pixels

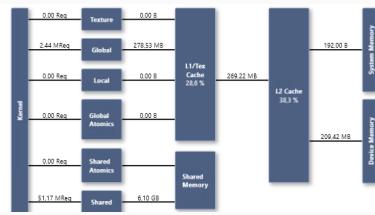
```
// TILE_WIDTH = blockDim.x + 2
__shared__ int tile[TILE_WIDTH][TILE_WIDTH]; // Alloc with the size of the block + border

int* block_ptr = in + ...;
for (int i = threadIdx.y; i < TILE_WIDTH; i += blockDim.y)
 for (int j = threadIdx.x; j < TILE_WIDTH; j += blockDim.x)
 data[i][j] = block_ptr[i * pitch + j];

__syncthreads();
```

## Stencil Pattern with Tiling Performance

- Global memory : 163 ms for a 24 MPix image
- Local memory : 116 ms for a 24 MPix image (30% speed-up)



47

48