

CS4803DGC Design and Programming of Game Consoles

Spring 2011

Prof. Hyesoon Kim



**Georgia
Tech**



College of
Computing



Lab #2


- Floating point multiply and add operation:
 - 2 FP operations
- Please look at PTX instructions
- You might not get what the device query says: explain why...
- Objdump will provide more precise results but for this assignment, just use ptx.
- Arithmetic Intensity: $\text{math operations per memory op} = \frac{\text{Sum of FP operations}}{\text{Sum of \# of transferred bytes}}$



Register Read & ILP & TLP

- Register read is fully pipelined.
- Back-to-back operation is in the critical path
- ILP across warps (~= TLP) can hide the latency of back-to-back

R1= R2+R3
R4= R1+R4



1 warp
24 cycles delay between 2 insts

1 warp
24 cycle delay is hidden by TLP



So...

```
loop{  
    a = a+c;  
}
```

dependent instructions across loops



Memory Operations

```
for (ii=0; ii<2000; ++ii) {  
  ref=base + tx;  
  sh_ref=base+tx;  
  temp[sh_ref] = dm[ref];  
}
```


```
for (ii=0; ii<2000; ++ii) {  
  ref=base + (16*ii)+tx;  
  sh_ref=base+(16*ii)+tx;  
  temp[sh_ref] = dm[ref];  
}
```

- Any performance difference?
- DRAM row buffer hit and miss will make a big difference




Memory access patterns

- coalescing




- Uncoalescing (Braid's lab)





Uncoalesced



- Mem addr = (tid)*X+Y + ii (loop iteration)
- And vary X and Y to generate different access patterns




DETOUR: DRAM




Hardware Structures

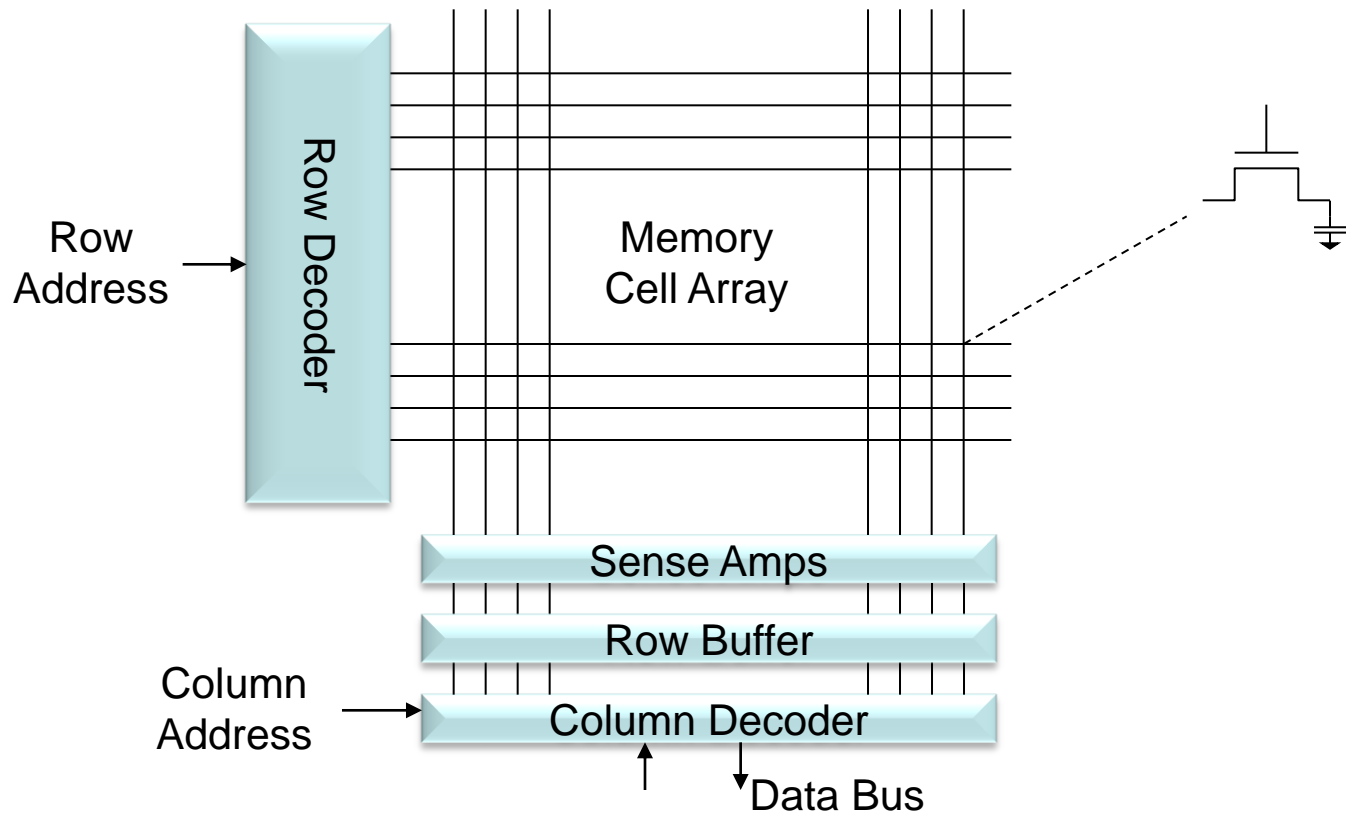
SRAM



DRAM

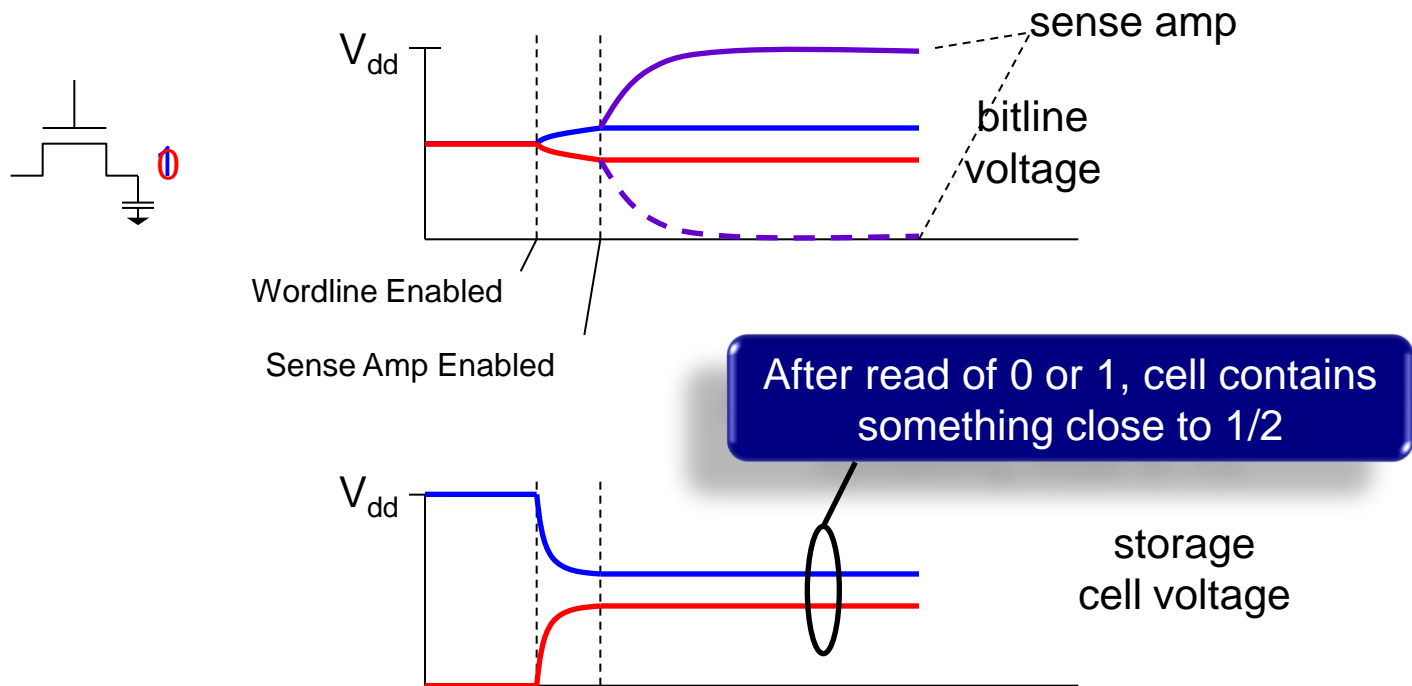


DRAM Chip Organization





Destructive Read



DRAM refresh is necessary to keep the data as well



DRAM Row Buffer

- Row buffer hit and miss penalty
- CAS+RAS+Precharge
- CAS
- Bank conflicts
- DRAM access time varies 10x



Ann.

- Lab #2: 7% → 10%.
- Friday 6 pm: Extra 10%.
- Extended due: Monday 6 pm
- One more pole for make-up class.
- Newsgroup participation will provide bonus points opportunities.