


# Performance optimization

*Bootcamp for SahasraT  
8th September 2018*

*Aditya Krishna Swamy  
[adityaks@iisc.ac.in](mailto:adityaks@iisc.ac.in)  
SERC, IISc*

# What is performance optimization

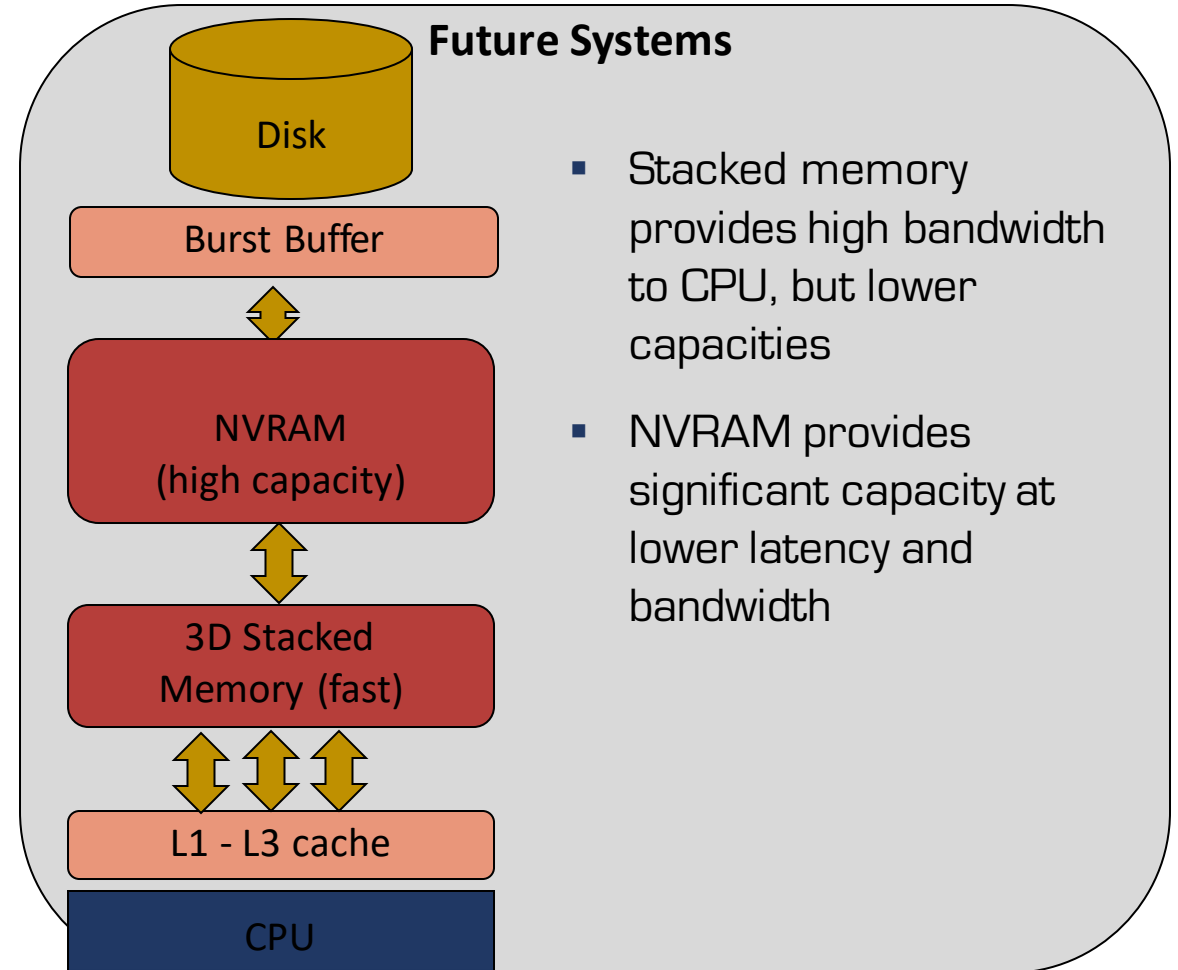
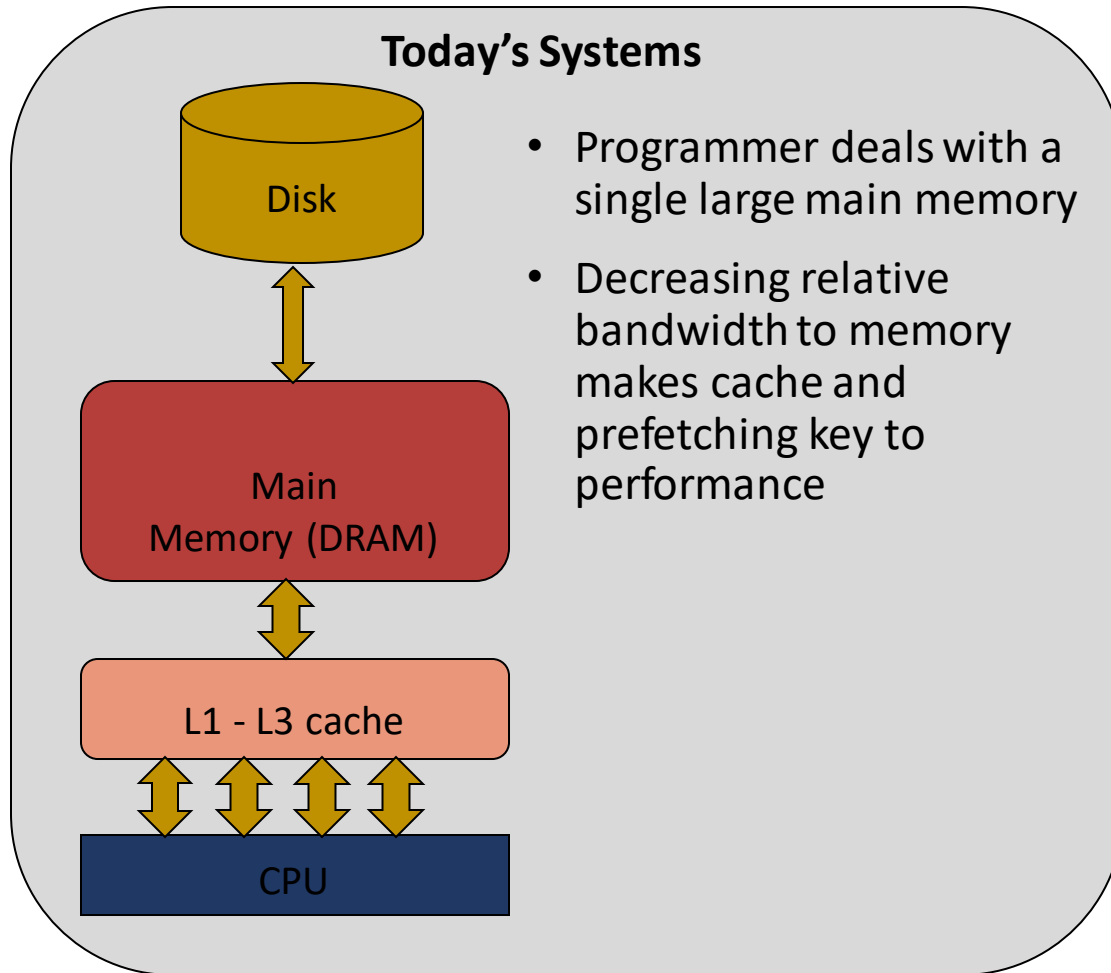
- A working code execution consumes resources – CPU, memory, power...
- Optimization - efficient use of these resources for faster results
- Available resources increasing rapidly



Scientific computing requires broad expertise  
*in addition to the research domain*


Knowledge of computer architectures, mathematical models  
and numerical algorithms

# Memory systems are getting increasingly Complex



Future Challenges: Treat stacked memory as cache, or explicitly managed?  
How can we take advantage of non-volatile nature of main memory?

# The reason for disruption $10^{15} \rightarrow 10^{18}$ FLOPS FLOPS

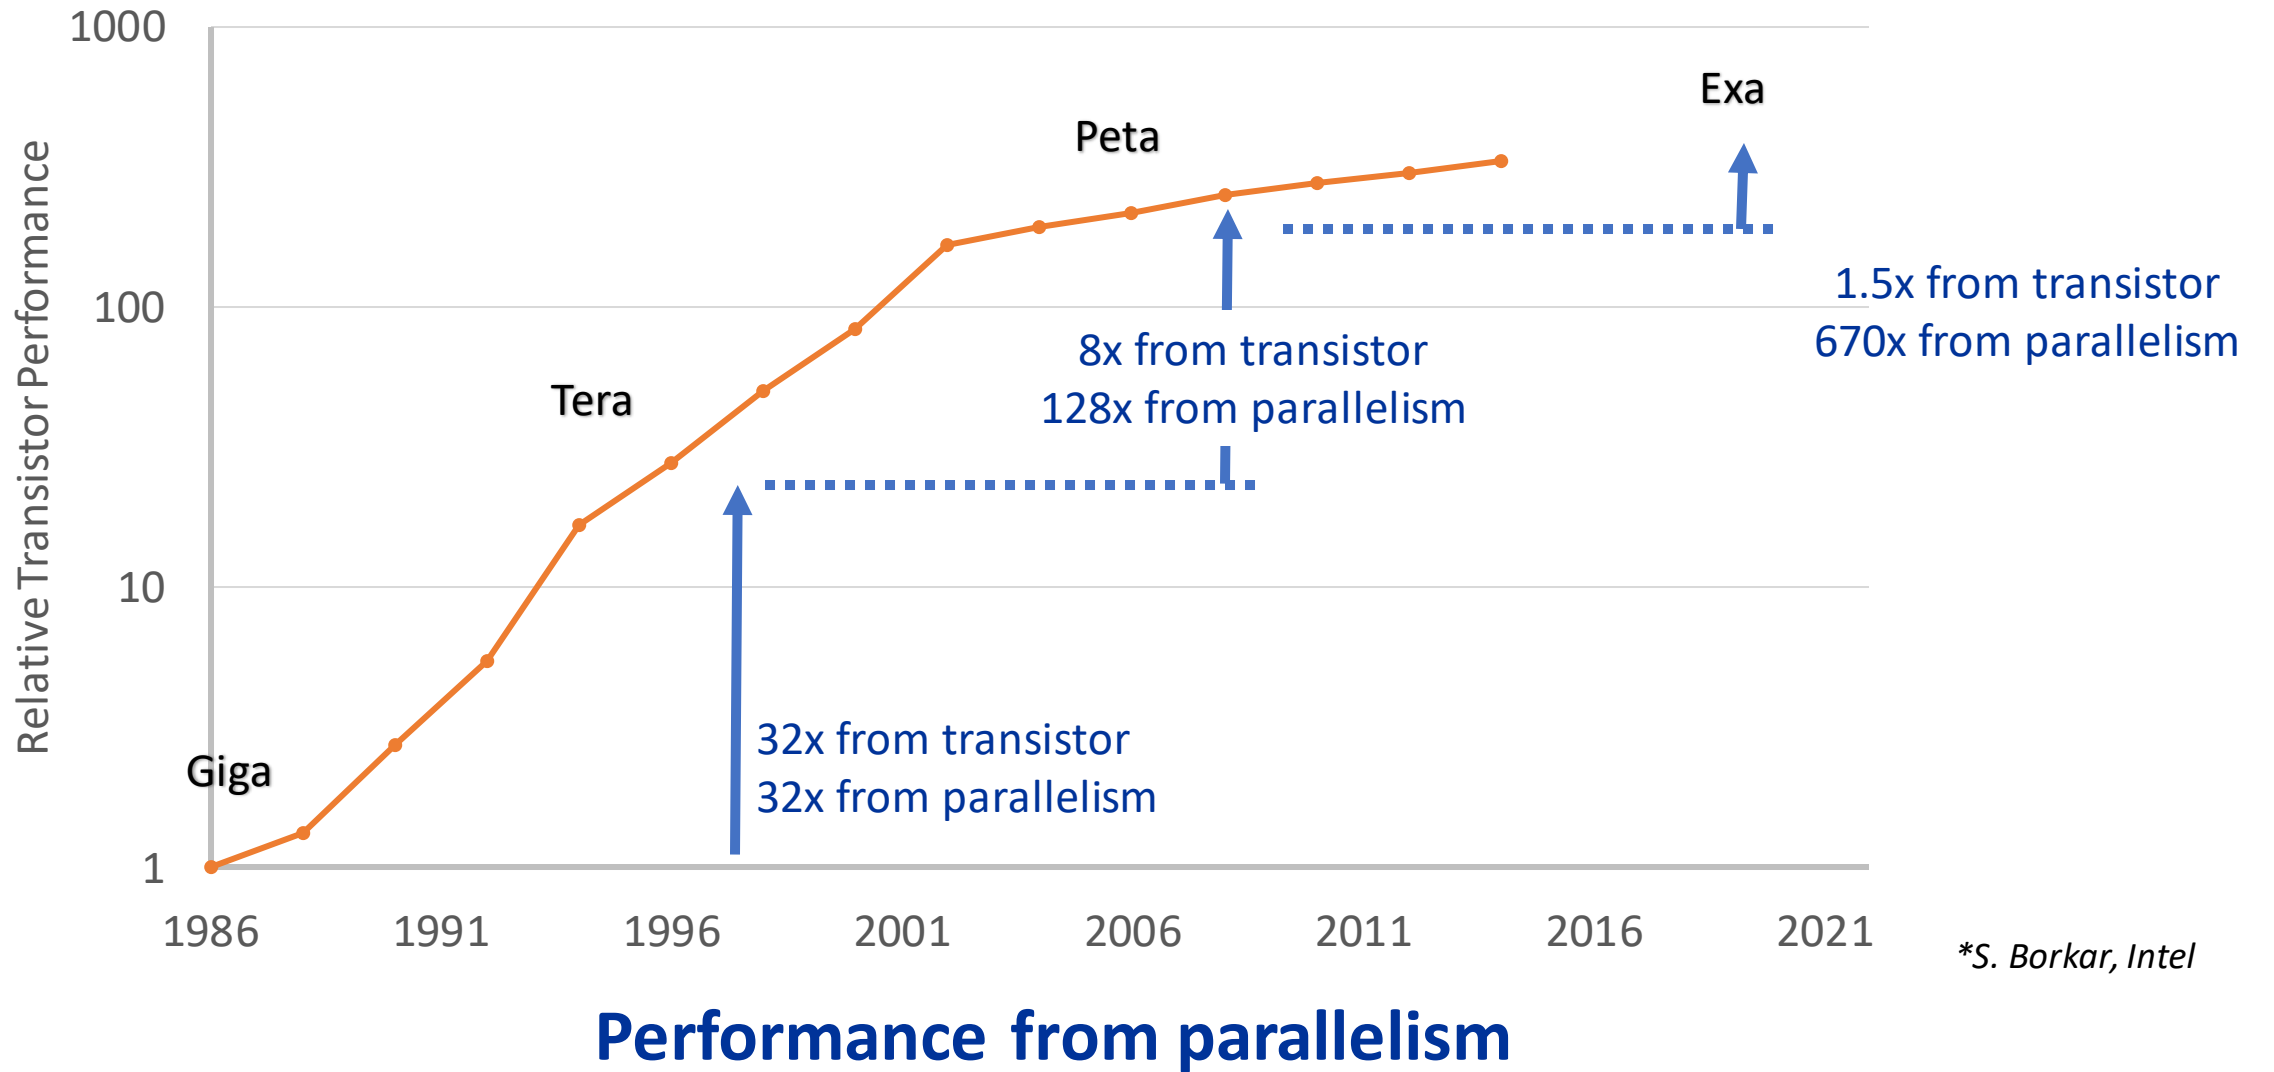


Moore's Law still holds  
Number of transistors  
continues to increase

Dennard scaling does not:  
Performance/Clock has flattened  
out because  
at 90nm transistor gates became  
too thin to prevent current from  
leaking out into the substrate

ILP: Instruction-Level Parallelism

# From Giga to Exa, via Tera & Peta\*



\*S. Borkar, Intel

# Practical example