MARS: Adaptive Remote Execution Scheduler for Multithreaded Mobile Devices

*Asaf Cidon**, Tomer M. London*, Sachin Katti, Christos Kozyrakis, Mendel Rosenblum

Stanford University

*Equal contributors

New Class of Mobile Applications

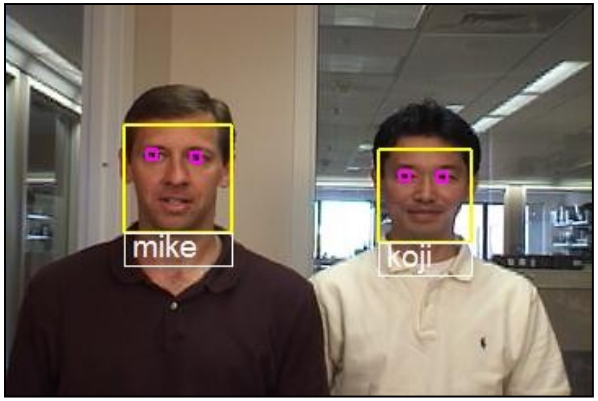

Computer Vision

Augmented Reality

Motion Sensing

Mobile Client Trends

- Mobile CPU performance increasing
  - Hitting ‘energy wall’
- Can we improve performance and reduce energy consumption?
- Opportunity: network bandwidth increase→ utilize the cloud

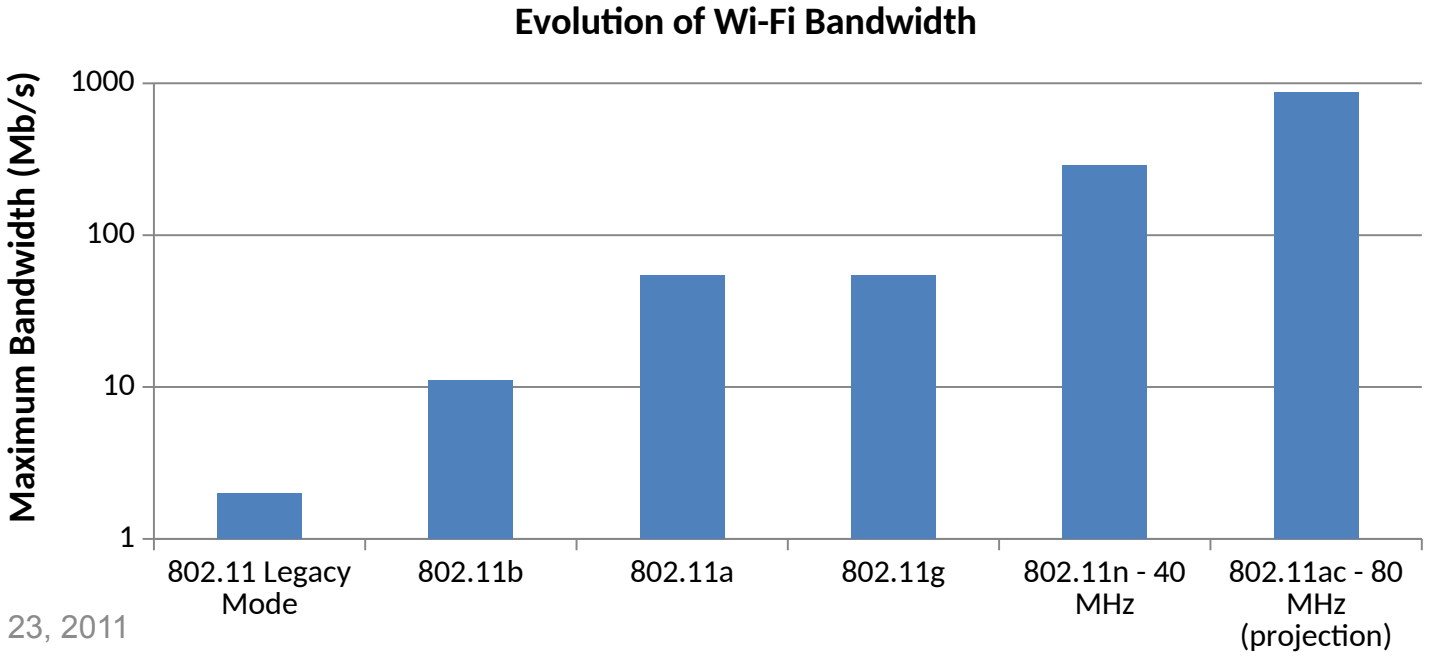

Static Client-Server Partitioning Doesn't Work

- Dynamic resources:
  - Network bandwidth and latency
  - Available CPU, memory
- Same code, different platforms:
  - Smartphones (single-core, multi-core)
  - Tablets

MARS: Adaptive Remote Execution

- Opportunistically offload computations to remote server
  - Enhance computational capabilities
  - Decrease energy consumption

- Make dynamic decisions
  - Adapt to network and CPU variability

Mobile Device

Data Center

Agenda

1. **Design of MARS**
2. Simulator Results and Analysis
3. Conclusions

Existing Remote Execution Systems

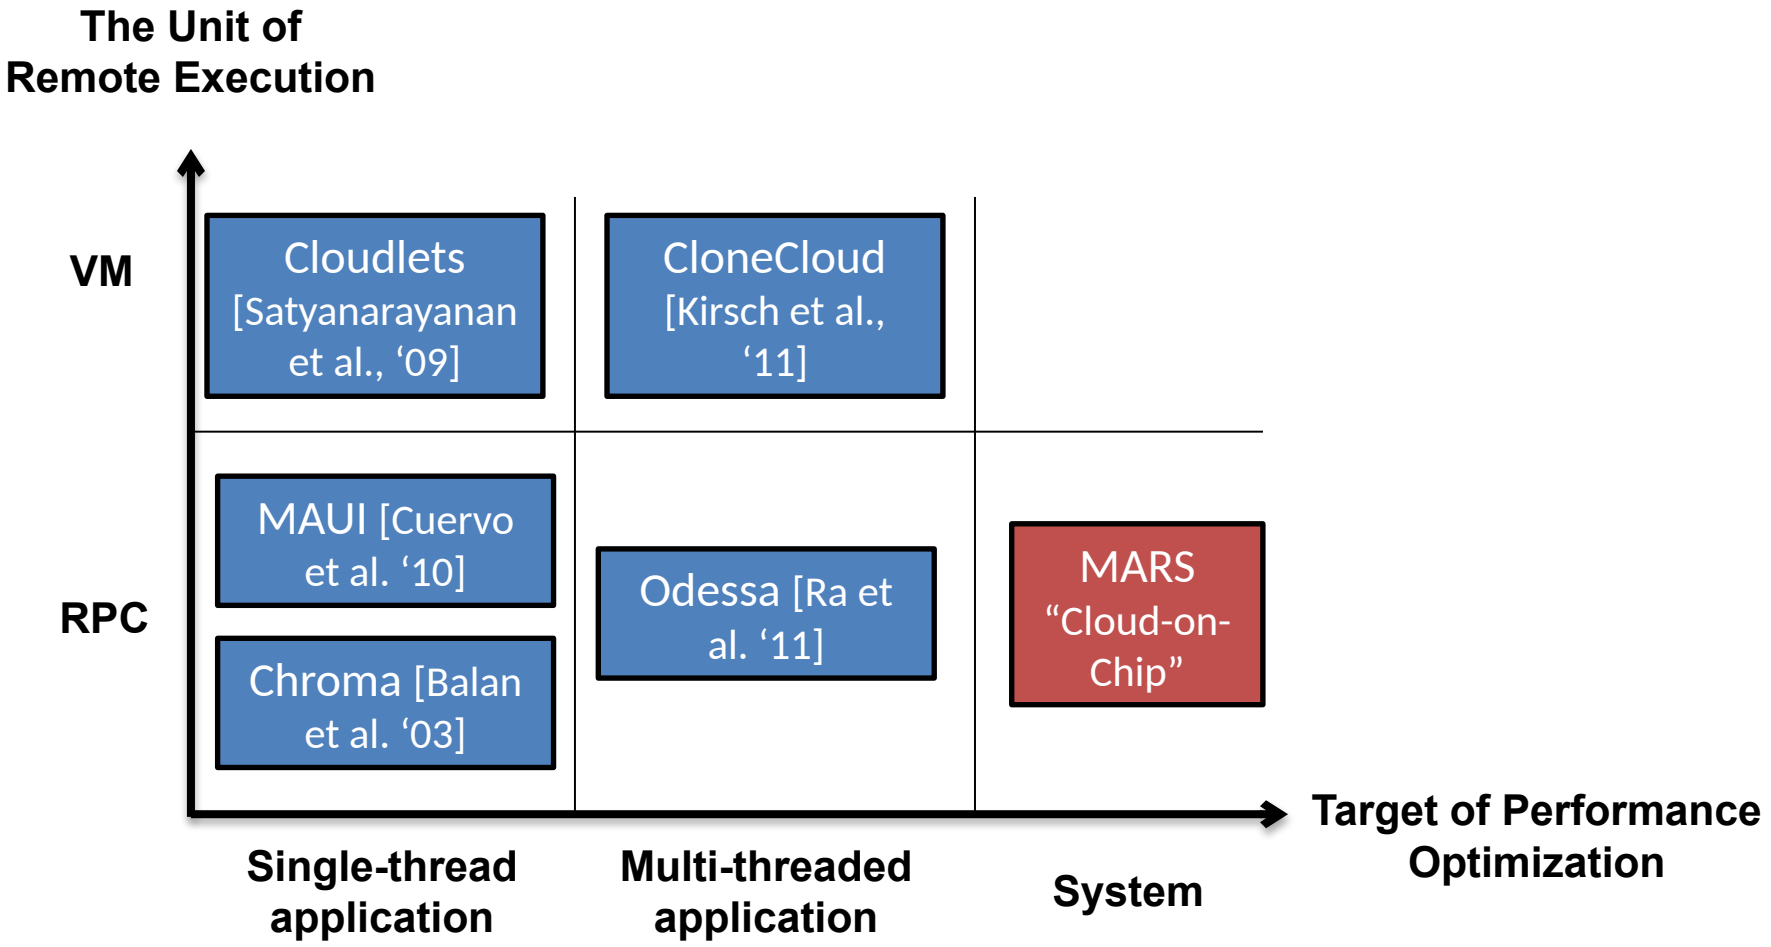

Previous Systems: Application Partitioning

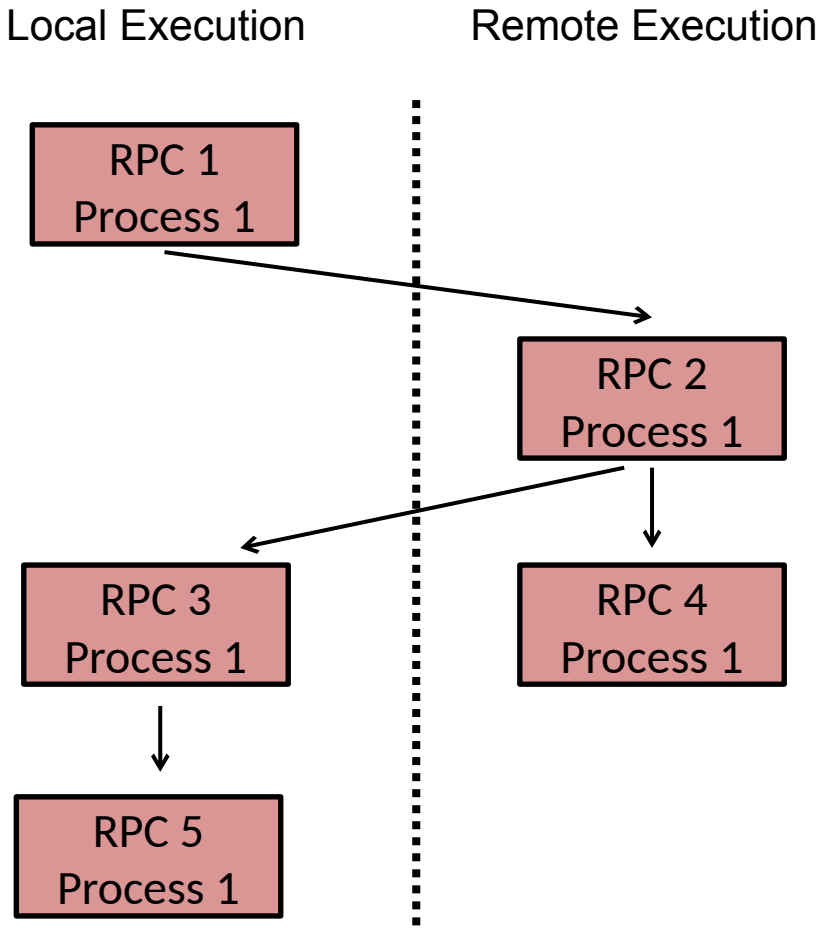

MARS "Cloud-on-Chip": System Scheduling

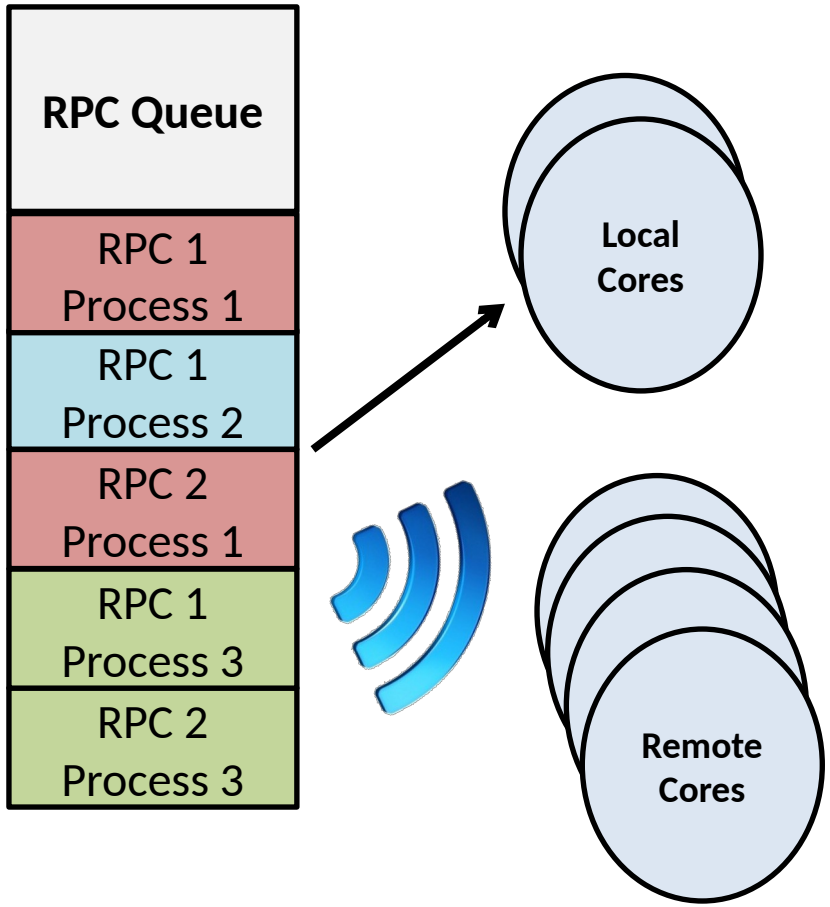

Greedy Algorithm

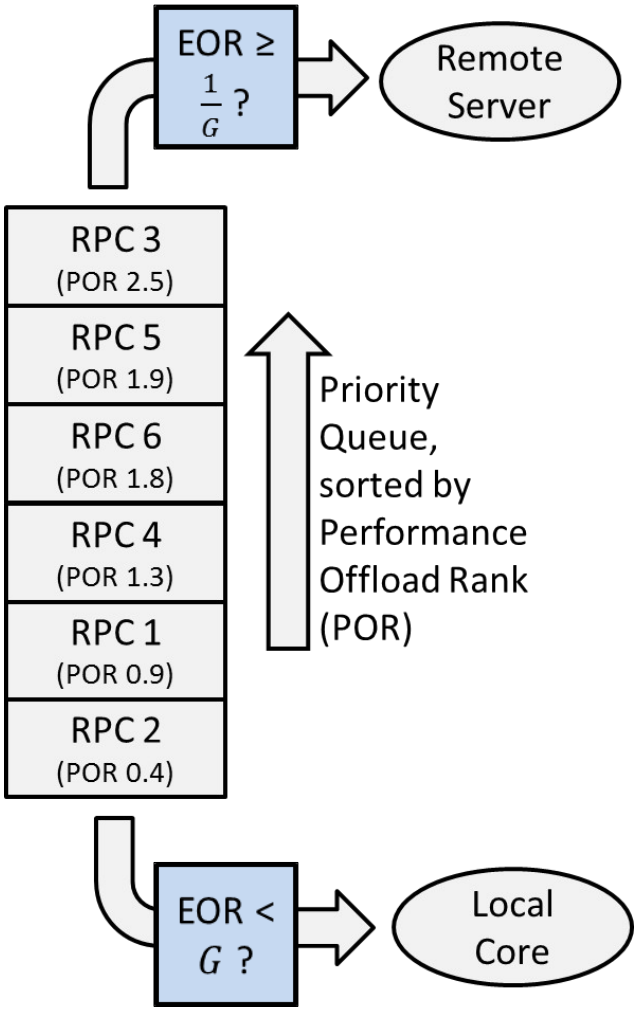

$$POR(RPC) = \frac{LocalExTime(RPC)}{RemoteExTime(RPC) + NetDelay(RPC)}$$

Higher POR: better **performance** gain from offloading

$$EOR(RPC) = \frac{LocalExTime(RPC) \cdot LocalPower}{NetworkEnergy(RPC)}$$

Higher EOR: better **energy** saving from offloading

Controller Algorithm

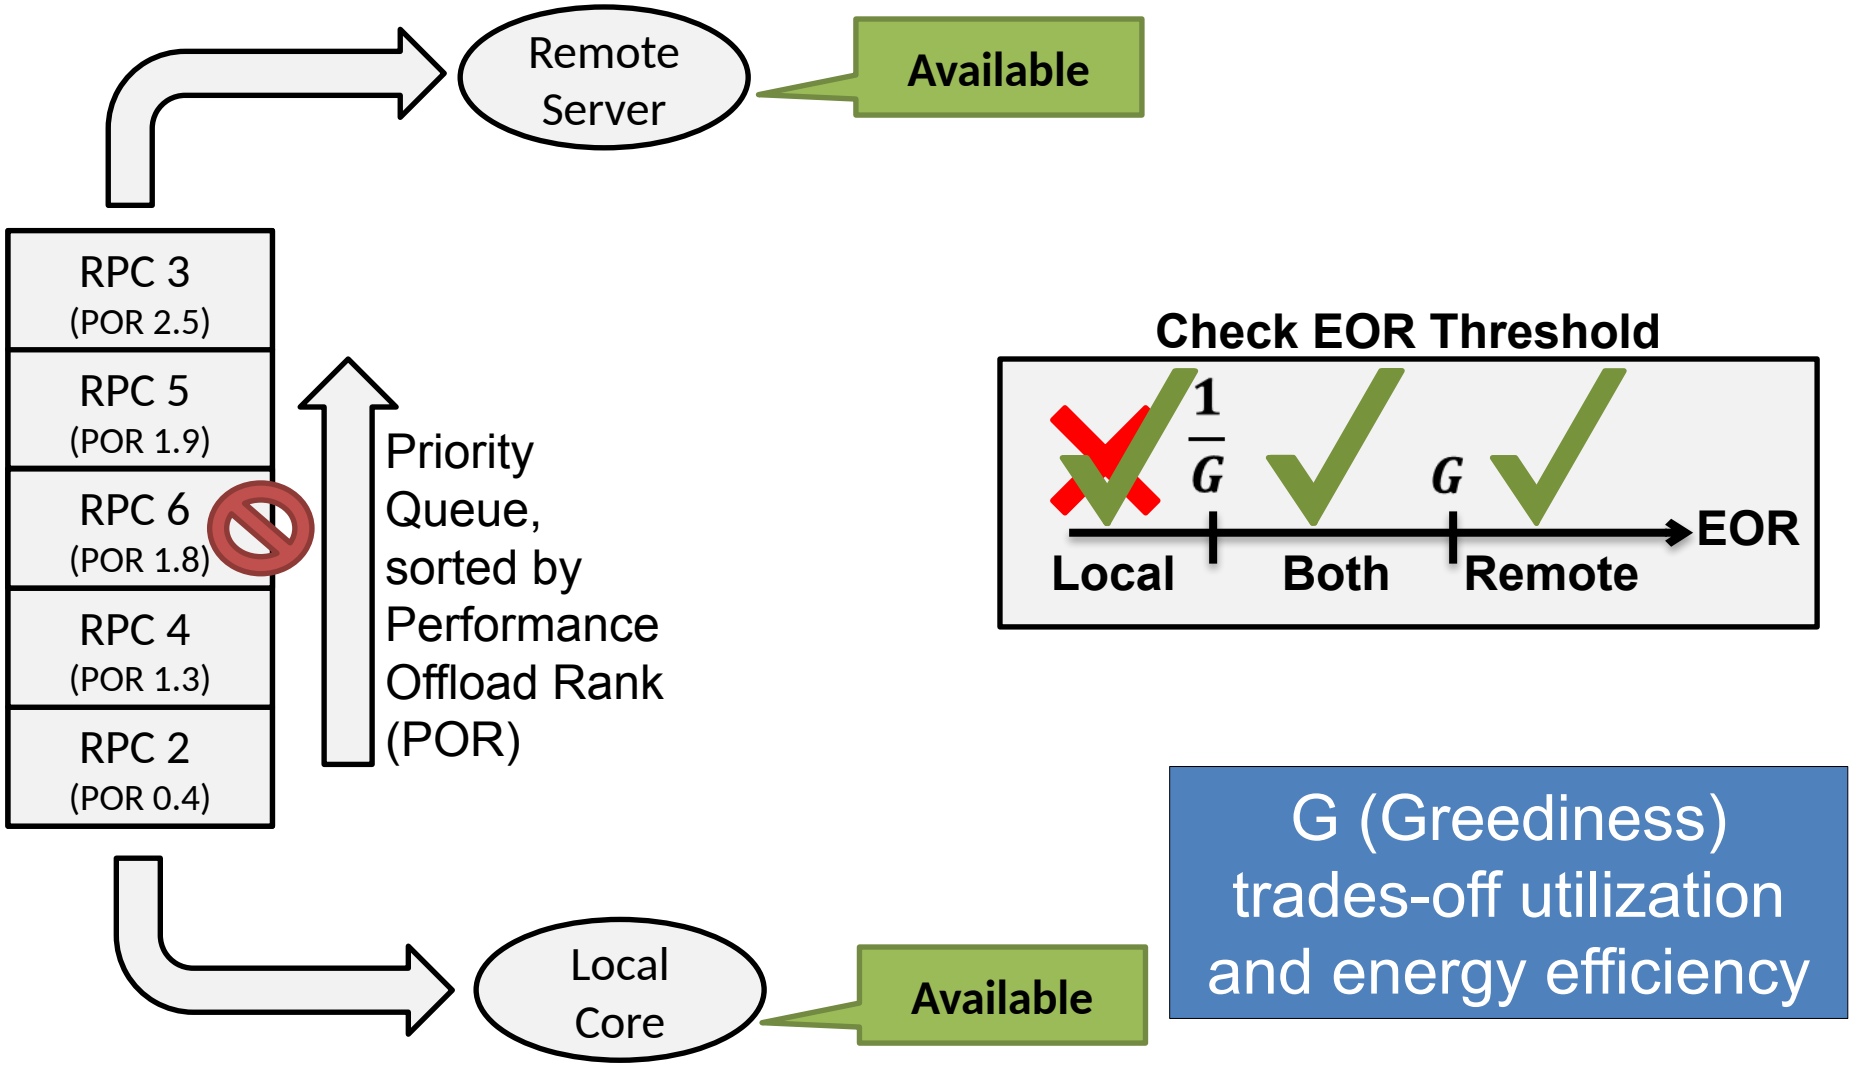

Agenda

1. Design of MARS
2. **Simulator Results and Analysis**
3. Conclusions

Remote Execution Applications

Augmented Reality

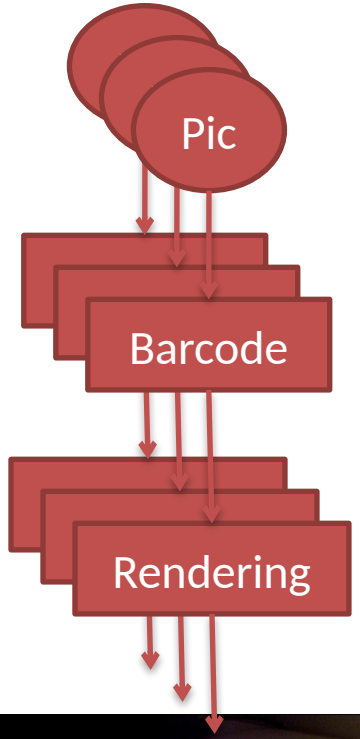

Face Recognition

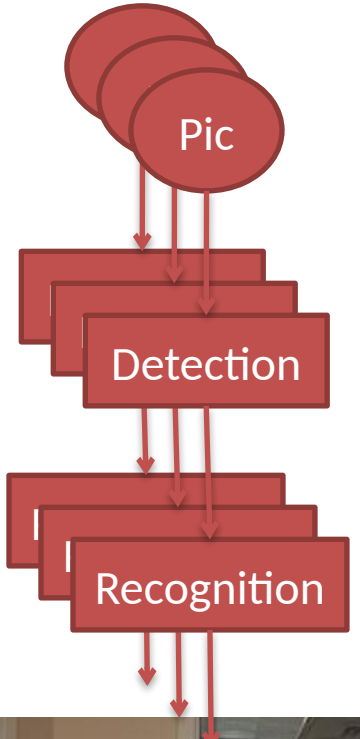

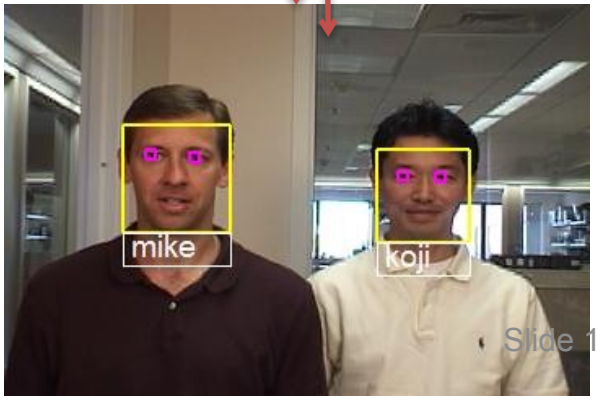

Simulator Methodology

- Trace-driven simulation
- Clients:
  - Nokia N900 (single core)
  - NVIDIA Tegra 250 (multicore)
- Server:
  - Amazon EC2 Opteron 2007
- Networks:
  - Outdoors Wi-Fi
  - Indoors Wi-Fi
  - 3G

MARS vs. Static Policies

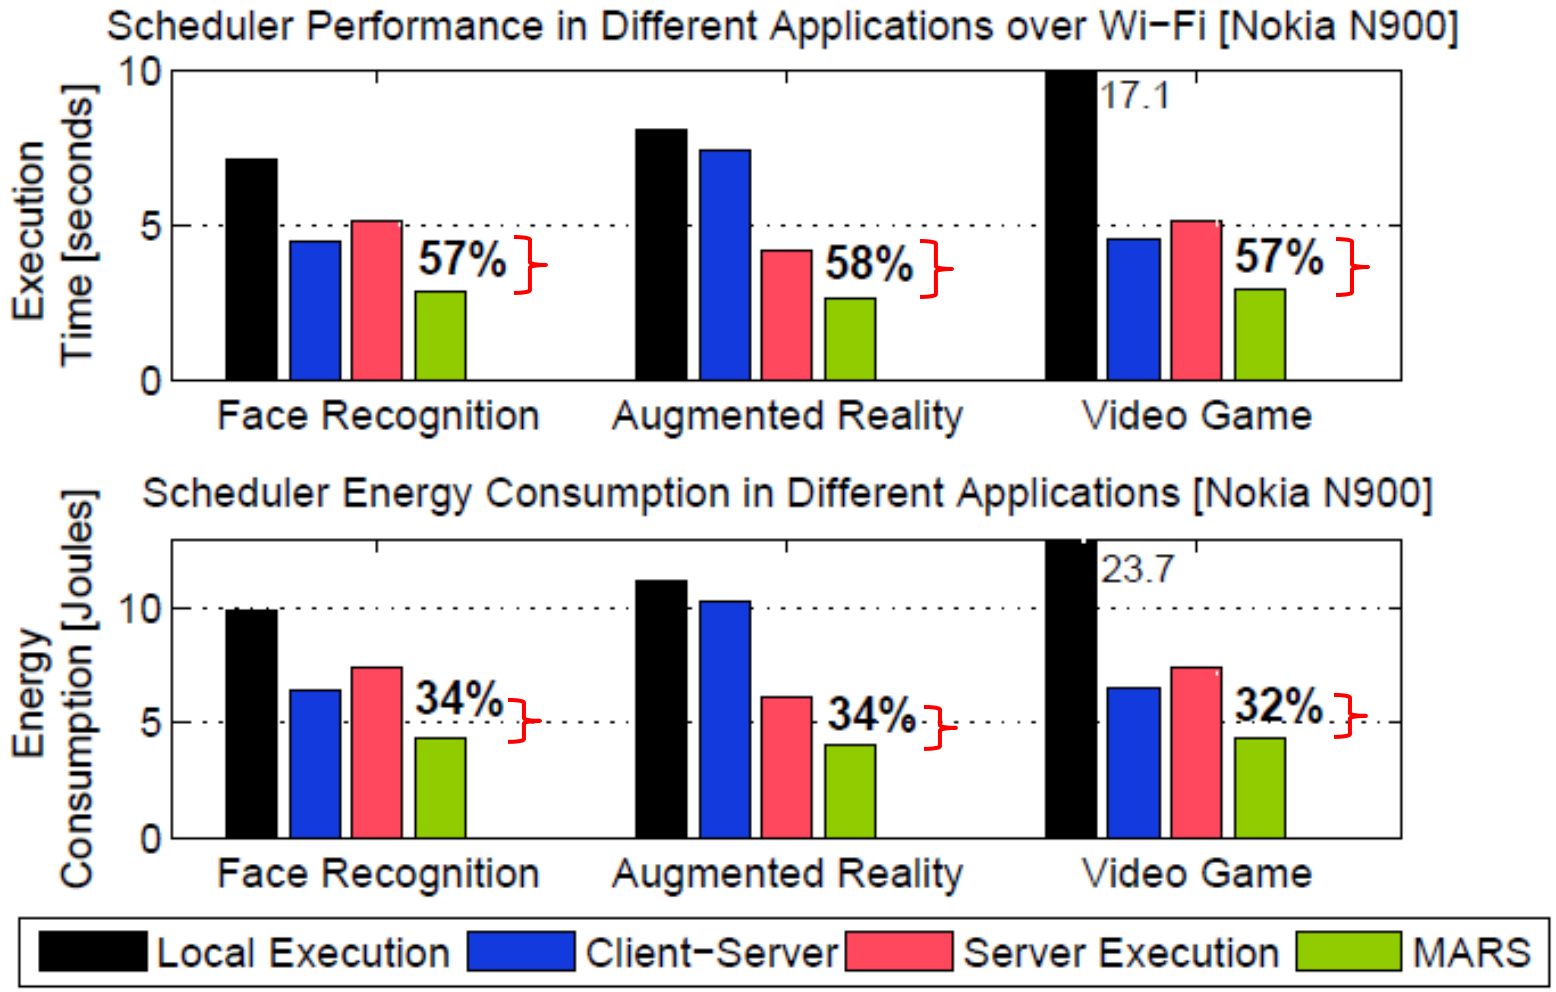

Nokia N900 Power Consumption

	Wi-Fi	3G
Idle Network Power	1.31 Watts	0.66 Watts
Upload Network Power	1.464 Watts	2.36 Watts
Download Network Power	1.39 Watts	2.26 Watts
Upload Network Power Overhead	**10.51%**	**72.03%**

- WiFi: Performance and energy are highly correlated
- 3G: trade-off performance and energy

Same Application, Different Networks

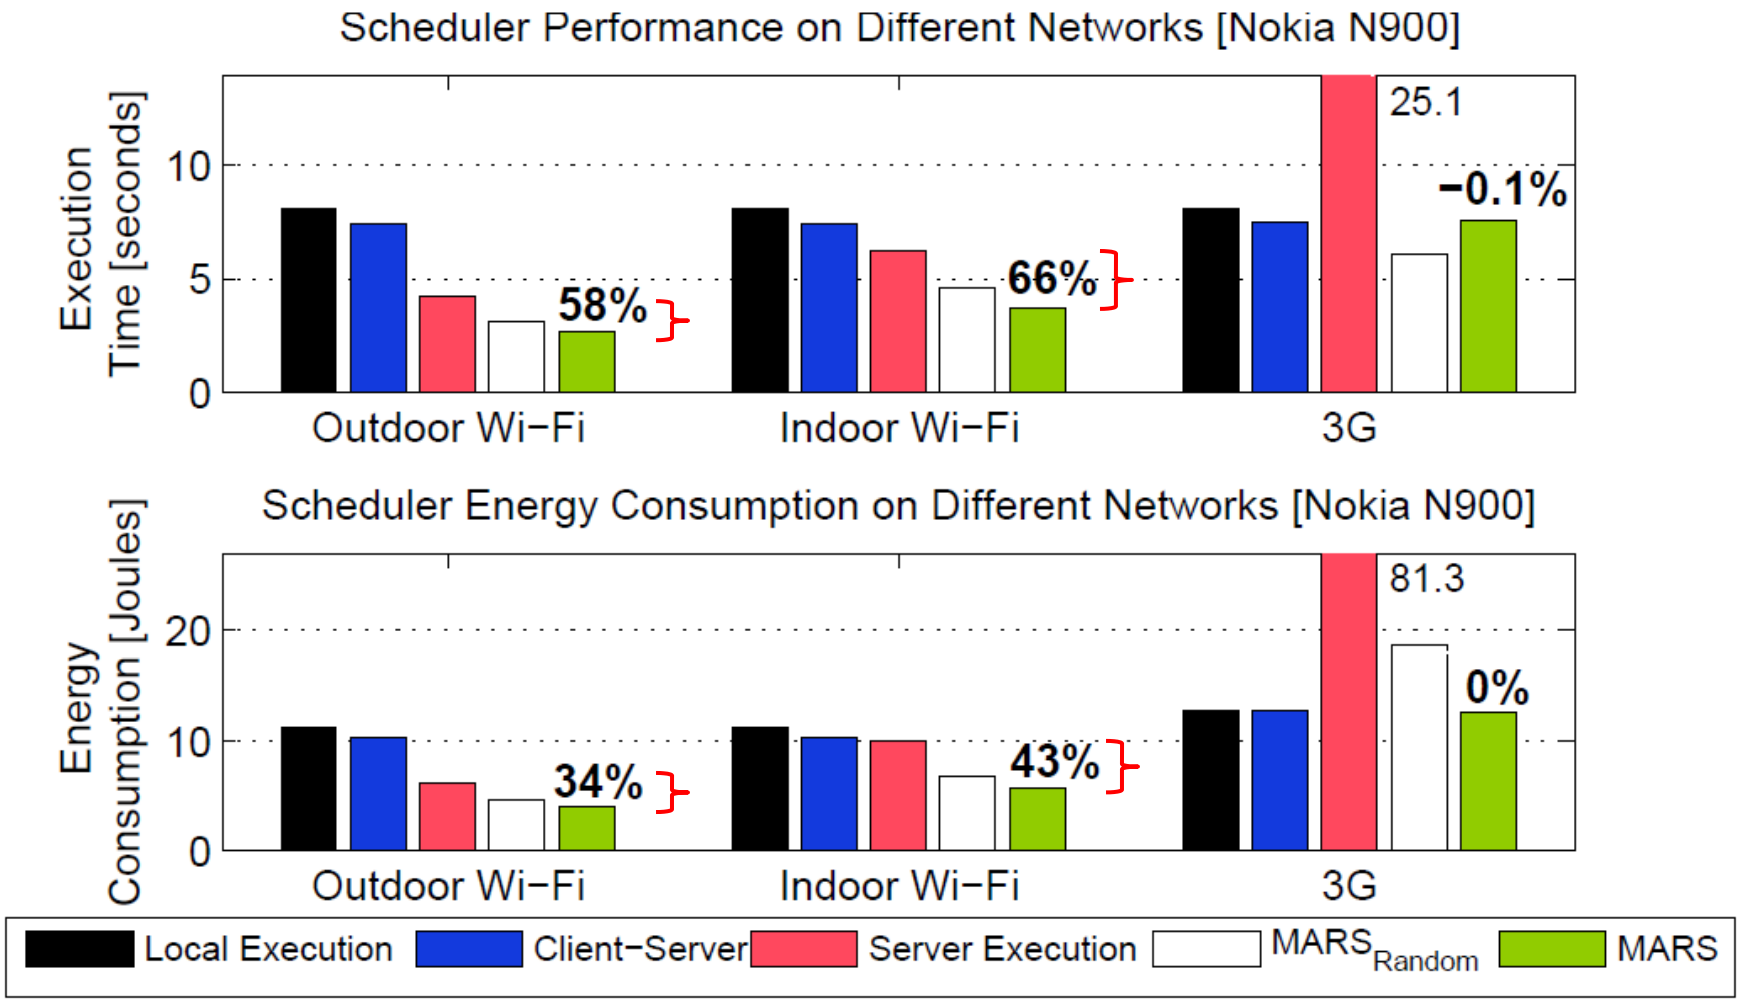

Remote Execution with Multicore

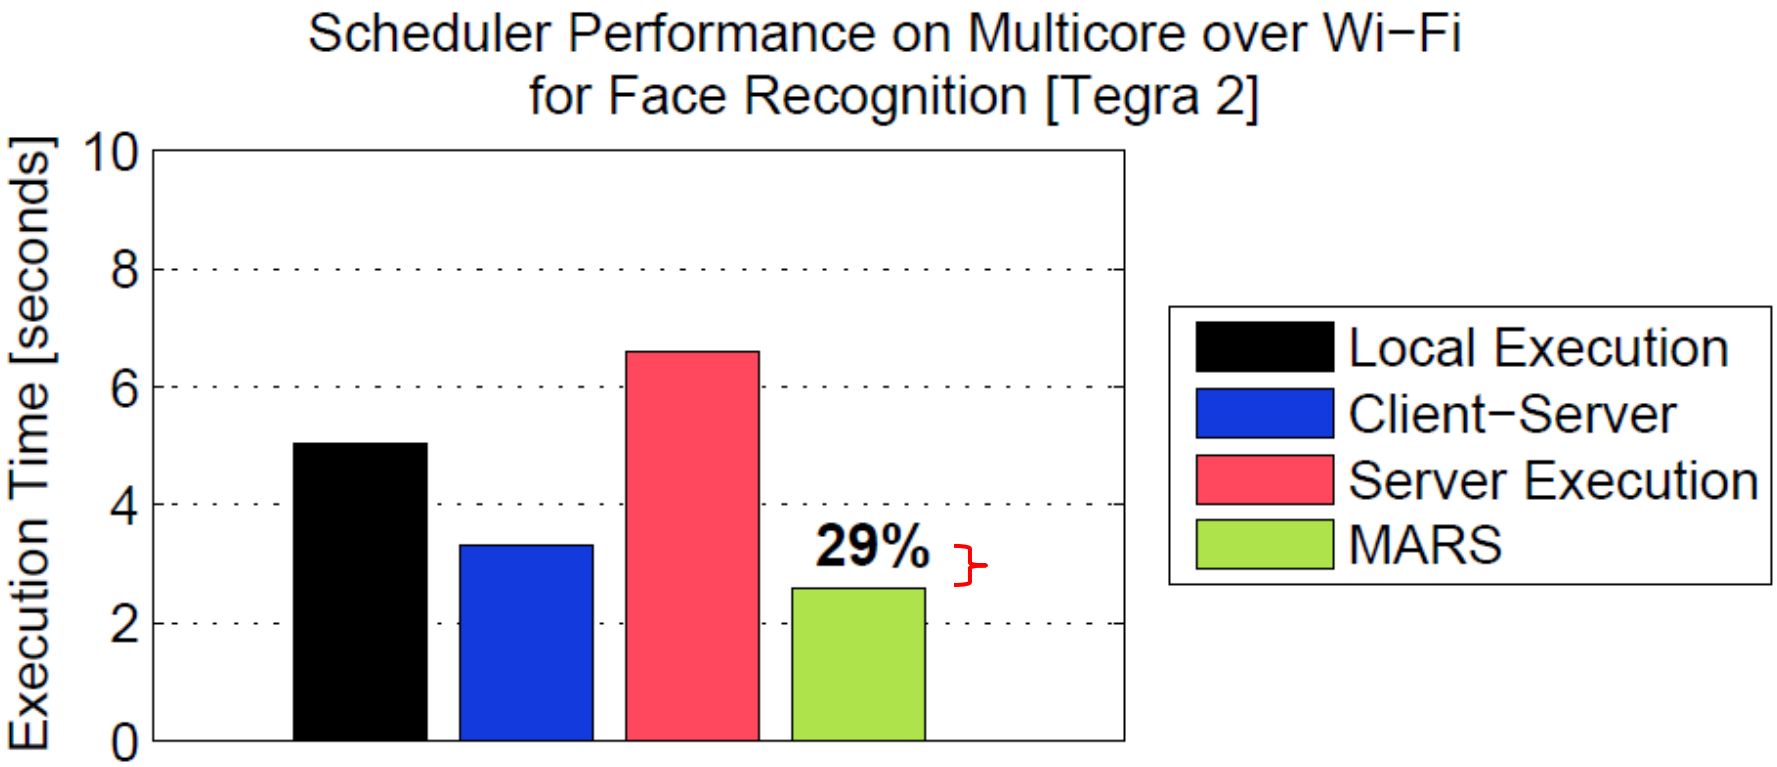

Agenda

1. Design of MARS
2. Simulator Results and Analysis
3. **Conclusions**

Conclusions

1. Can't always be greedy
   - Performance and energy trade-off
2. MARS is optimized for multiple parallel applications and cores
3. MARS "Cloud-on-Chip": validation of system-level remote execution scheduling
   - 57% performance increase, 33% energy savings