Apollo: Adaptive power optimization and Control for the land Warrior

Massoud Pedram
Dept. of EE-Systems
University of Southern California
Niraj K. Jha
Dept. of EE
Princeton University

April 26, 2001

Outline

- Part I: Dynamic Power Management and Power-Aware Architecture Reorganization
  - OS-directed power management
  - Architecture organization techniques
  - Apollo Testbed
  - Summary I
- Part II: Power-Aware Behavioral and System Synthesis
  - Input space adaptive software
  - Power-aware dynamic scheduler
  - High-level software energy macromodels
  - Leakage power analysis and optimization
  - HW-SW co-synthesis with DR-FPGAs
  - Low power distributed system of SOCs
  - Summary II

Introduction

- Land Warrior (LW) system's design objectives include:
  - Situational awareness
  - Power awareness
- Average power reduction of the LW system while meeting key minimum performance and quality-of-service criteria is an important design driver
- Much of the power savings is at the system-level and can be achieved thru dynamic power management (DPM) and power-aware architecture organization
- Our research focuses on these two approaches to system-level power reduction and is expected to deliver 2-4 X power savings for the LW system

Land Warrior

Project Overview

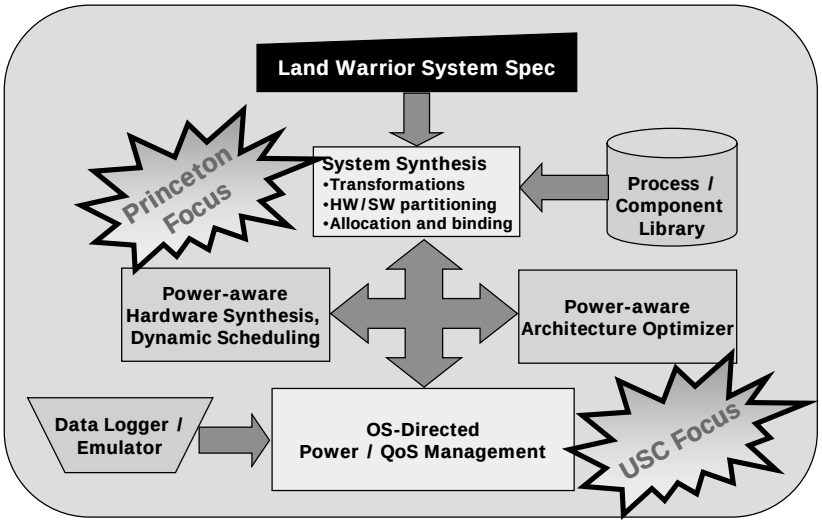

OS-directed Power Management

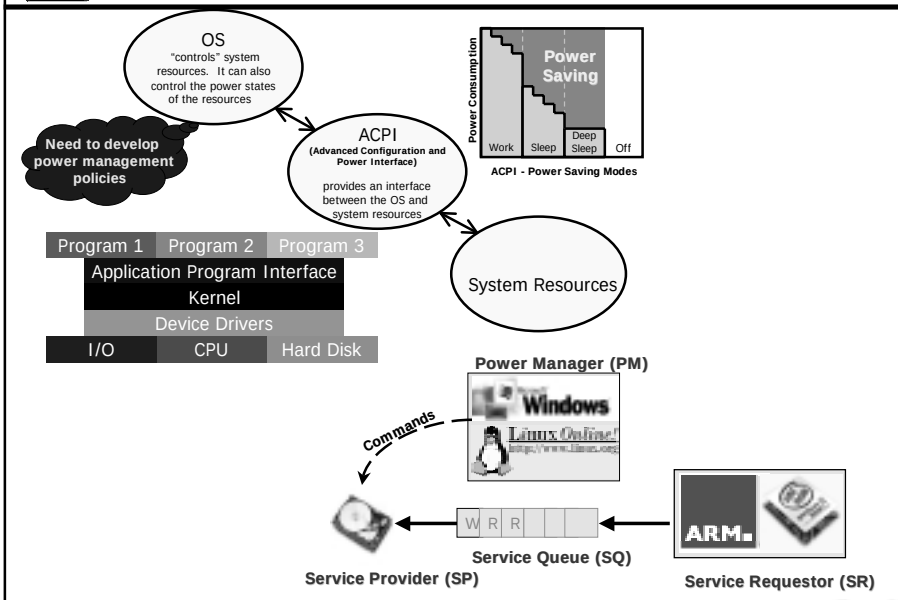

Policy Optimization Flow

- Power management policy optimization flow:
  - Build stochastic models of the service requesters (SRs) and service providers (SPs)
  - Construct a continuous-time Markovian decision process model (CTMDP) of the power-managed system
  - Obtain stochastic parameter values by static profiling and runtime monitoring
  - Solve the resulting power optimization problem under performance and quality-of-service constraints to obtain the optimal policy

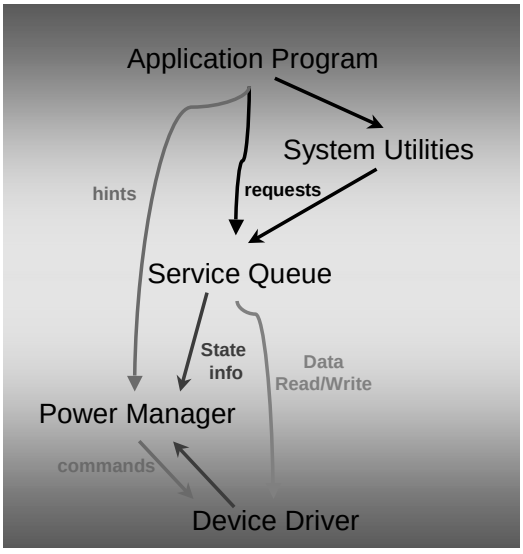

Example: A Fujitsu Hard Disk

- Two abstract models of the Fujitsu Hard Disk

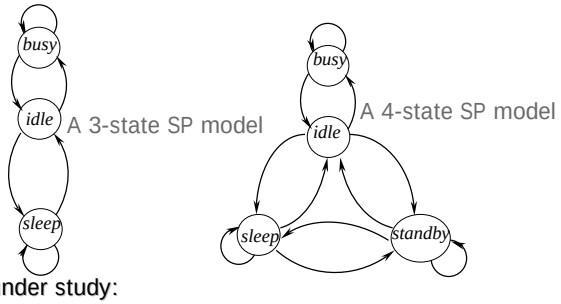

- Policies under study:
  - "Always On" policy
  - "Greedy" reactive policy
  - "Time Out" policies with different time-out values
  - "N" policies with different N values
  - 3CTMDP: CTMDP policies using the 3-state SP model
  - 4CTMDP: CTMDP policies using the 4-state SP model
  - CTMDP-Poll: 3CTMDP with polling to address the long tail of distribution

Experimental Results

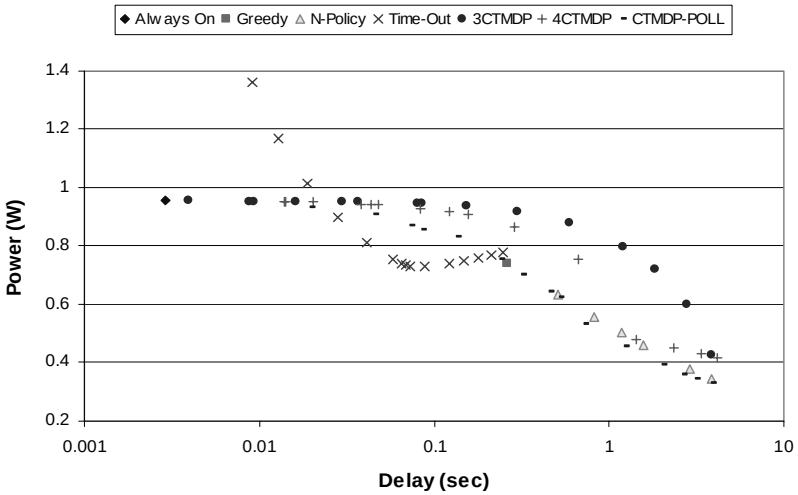

Power-delay trade-off curves for Uniform transition time distribution for the hard disk and a combination of Exponential and Pareto distributions for the request inter-arrival times

A Two-layered PM Strategy for the LW

Mode-based Decision Tree

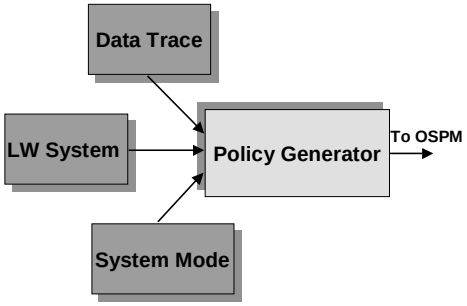

Parametrizable PQ policy

Apollo

On-chip Memory Bus Power Reduction

The processor-memory bus is highly capacitive. Significant power is consumed to drive these busses

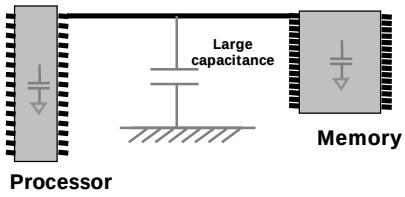

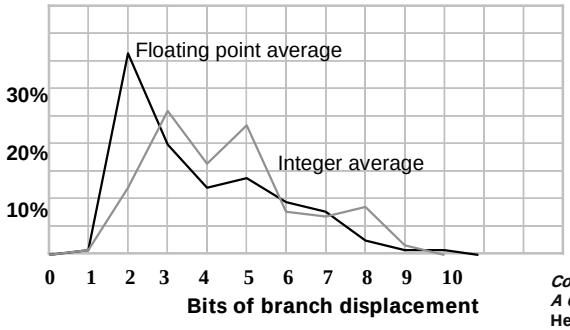

Computer Architecture, A Quantitative Approach, Hennessy and Patterson

Apollo

Address Bus Encoding: Alborz Code

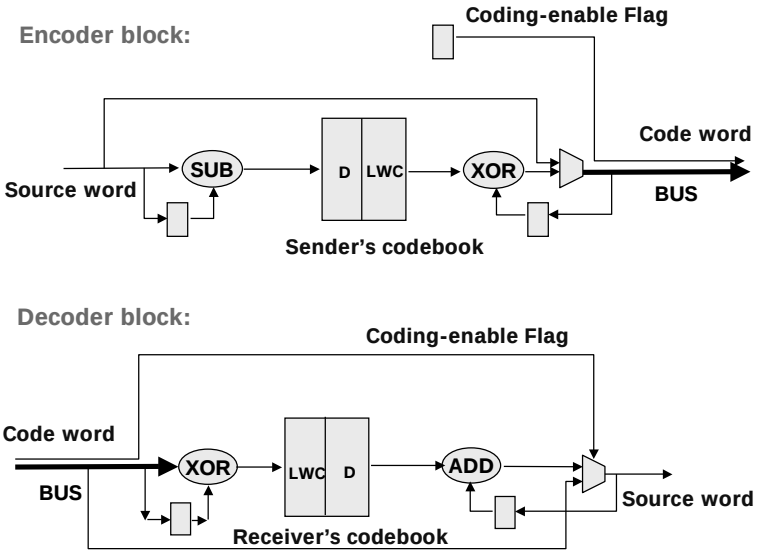

M. Pedram and N. Jha

Apollo

Experimental Results

Benchmark	Base Case Pwr	Fixed ALBORZ 512 Entry Redundant Pwr	Adaptive ALBORZ Redun.		Adaptive ALBORZ 32 Entry Irredun. Pwr
			32 Entry	64 Entry	
Art	34.3	2.765	2.466	2.110	2.382
		92.0%	92.9%	93.9%	93.1%
Gzip	34.7	4.005	3.103	2.651	3.483
		88.5%	91.1%	92.4%	90.0%
Vortex	37.3	6.048	6.244	5.554	4.918
		85.9%	83.3%	85.2%	86.9%
Equak	35.4	6.383	7.673	6.472	6.576
		82.1%	78.4%	87.8%	81.5%
Gcc	35.4	6.373	7.651	6.881	7.490
		82.1%	78.5%	80.6%	78.9%
Parser	37.0	4.177	4.273	3.614	4.560
		88.8%	88.5%	90.3%	87.3%
Vpr	35.9	7.021	6.085	5.502	5.928
		80.5%	83.1%	84.7%	83.5%
% power reduction		85.3%	85.1%	87.0%	85.9%

M. Pedram and N. Jha

Apollo

Off-chip Memory Bus Power Reduction

- DRAM address bus is multiplexed between the row and column addresses

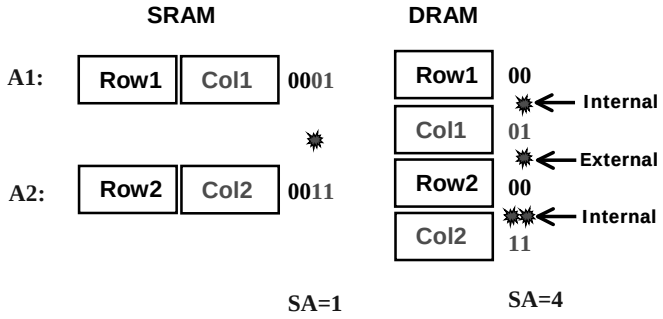

- Find a permutation (one-to-one function) that generates the minimum external switching activity

Apollo

Example

- Address space of 0, 1, 2, 3 = 2^2
- Multiplexed on a 1-bit bus
- 2^1 nodes and 2^2 edges

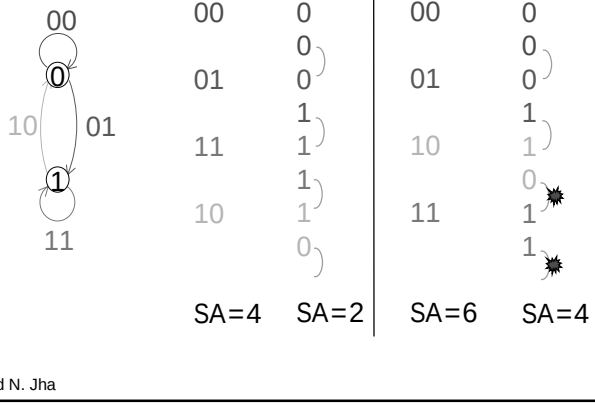

Apollo

RC Graph and Eulerian Cycle

- Construct a merged Row-Column Graph where nodes represent row and column addresses and edge (u,v) corresponds to binary number uv
- Find an Eulerian cycle on this graph using forward or backward traversal to obtain the Pyramid code

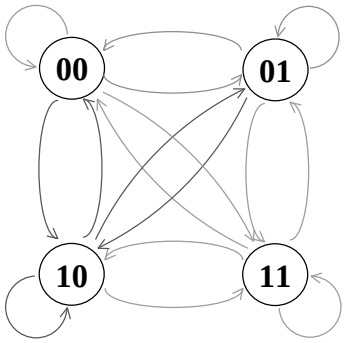

	00	01	10	11
00	0	8	11	1
01	10	9	14	6
10	15	13	12	4
11	7	5	3	2

Pyramid I – Forward Traversal

Experimental Results

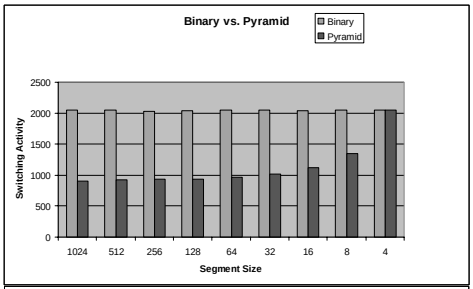

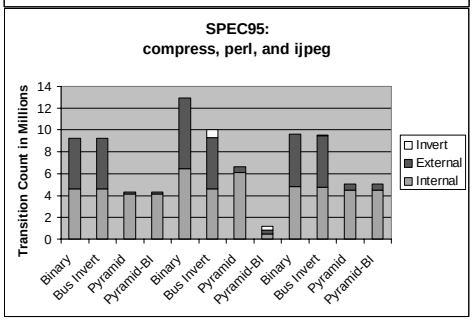

Apollo Testbed Block Diagram

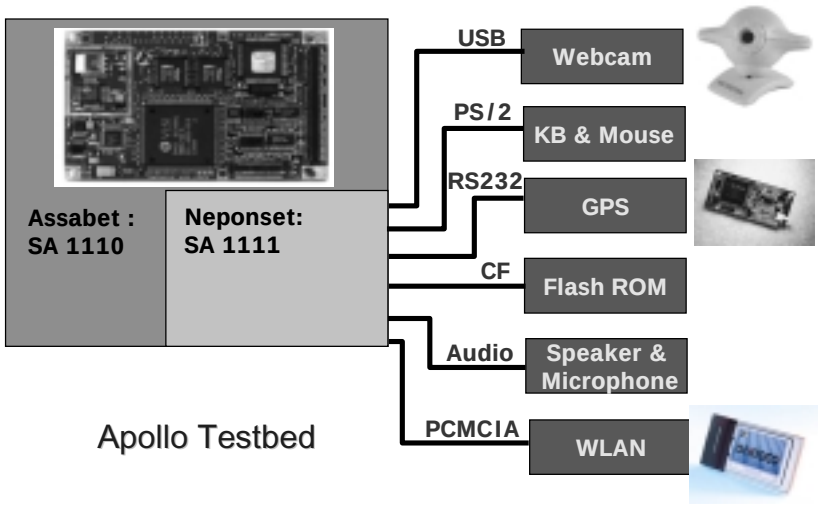

M. Pedram and N. Jha

Apollo

Device Information

- **Microtek EyeStar U2S Webcam**
  - USB interface
  - VISION CPiA chip (Color Processor interface ASIC)
  - 30 fps
  - 24-bit color
  - 640x480 pixels

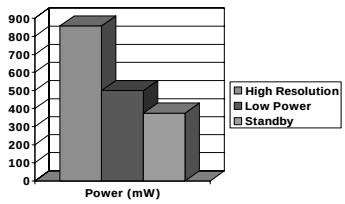

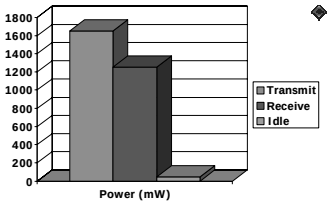

- **Cisco Aironet 340 Wireless LAN PC Card**
  - IEEE 802.11b
  - Data rate: 1, 2, 5.5 and 11 Mbps
  - Range: 120 m outdoors; 30 m indoors
  - Adjustable Transmit Power

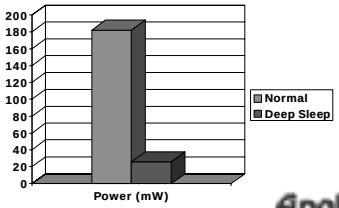

- **Trimble Lassen LP GPS Kit**
  - 8-channel, continuous tracking receiver
  - Update rate: 1 Hz
  - Position accuracy: 25 m
  - Velocity accuracy: 0.1 m/sec

M. Pedram and N. Jha

Apollo

Apollo Testbed Photo Shots

Hardware/Software Layers

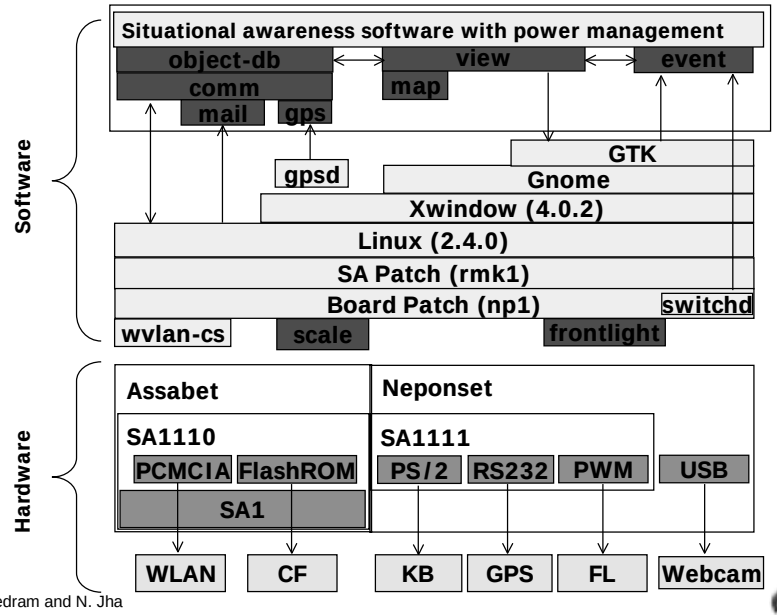

Map Viewer/Finder Snap Shots

M. Pedram and N. Jha

Apollo

Summary I

- The degree of power savings due to OSPM is dependent on the workload conditions and the minimum performance and quality-of-service requirements
- We expect a factor of 2-4 reduction in the total power dissipation of the LW system. This is being demonstrated on the Apollo Testbed
- Even higher power savings are possible if the processor, I/O device hardware, and device driver software are equipped with built-in power management facilities
- Future research work includes power reduction of the PCMCIA wireless LAN and the USB Webcam.
- The next generation of the Apollo Testbed will use the XScale processor from Intel. XScale microarchitecture offers a 4-8 X improvement in the power-performance metric over Intel StrongARM

M. Pedram and N. Jha

Apollo

Outline

- Part I: Dynamic Power Management and Power-Aware Architecture Reorganization
  - Introduction
  - OS-directed power management
  - Architecture organization techniques
  - Apollo Testbed
  - Conclusions
- Part II: Power-Aware Behavioral and System Synthesis
  - Input space adaptive software
  - Power-aware dynamic scheduler
  - High-level software energy macromodels
  - Leakage power analysis and optimization
  - HW-SW co-synthesis with DR-FPGAs
  - Low power distributed system of SOCs
  - Summary II

Input Space Adaptive Software Synthesis

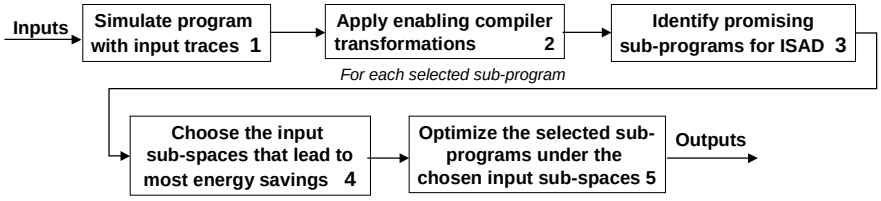

Improvements in performance and energy (*SPARClite*/*StrongARM*):

	Performance	Energy	Energy*delay
Maximum	7.2X/8.5X	7.6X/7.8X	54.7X/66.3X
Average	3.7X/3.0X	3.7X/3.1X	13.7X/9.3X

Dynamic Scheduling

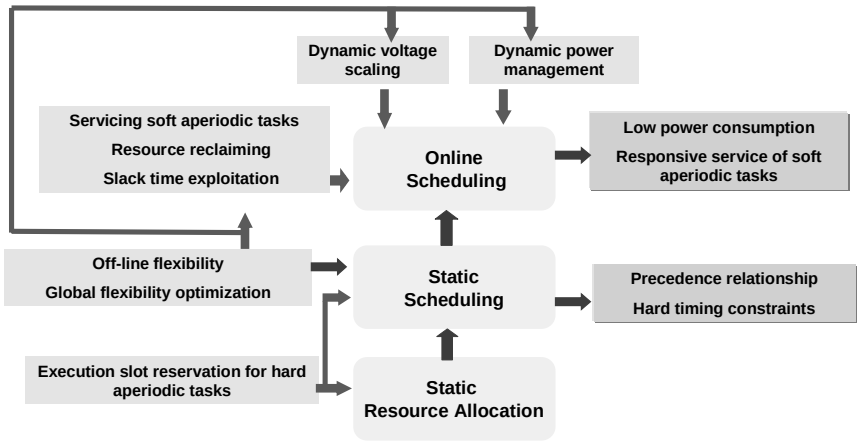

Power reduction over non-power-conscious scheme: ave.: 2.2X; max.: 3.1X

M. Pedram and N. Jha

Apollo

Battery-aware Scheduling

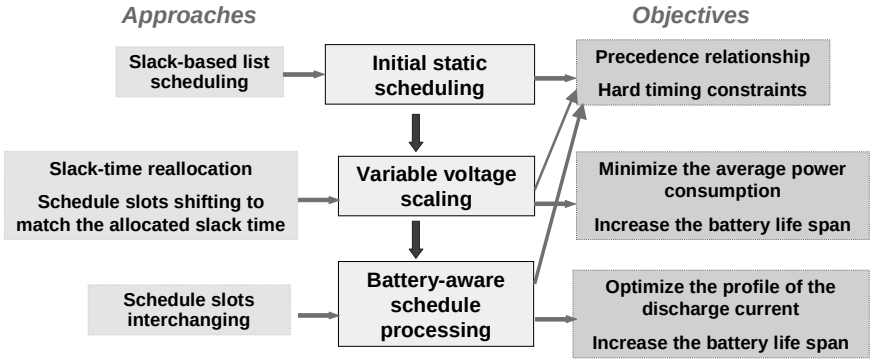

Battery lifespan extension: average: 1.3X, maximum: 1.6X

M. Pedram and N. Jha

Apollo

Energy Macromodels for Software

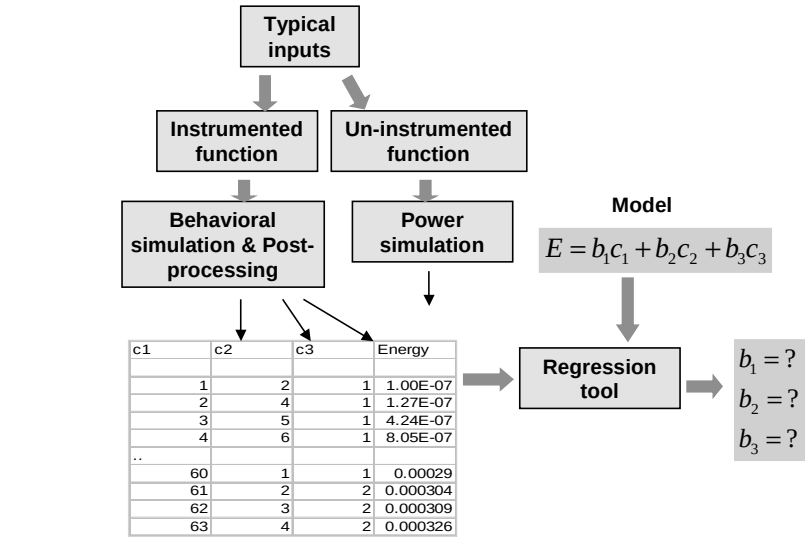

Complexity-based Model Results

Examples	Models	# of samples	SPARClite	
			Error	Speedup
chksum	$c_1 + c_2N$	400	1.4%	1361
igray	$c_1 + c_2log_2(N)$	2560	8.0%	540
edgedet	$c_1 + c_2M + c_3N + c_4MN$	1000	0.3%	673325
ins_sort	$c_1 + c_2N + c_3N^2$	250	6.7%	30050
mult	$c_1 + c_2L + c_3LM + c_4LMN$	625	2.4%	32213
myqsort	$c_1 + c_2N + c_3Nlog_2(N)$	250	5.3%	38155
msort	$c_1 + c_2N + c_3Nlog_2(N)$	1000	4.2%	126780
myfrag	$c_1 + c_2N$	1500	14.6%	81517

Ultra SPARC II (336 MHz) running SunOS 5.7

Real memory 4GB

Leakage Power Optimization

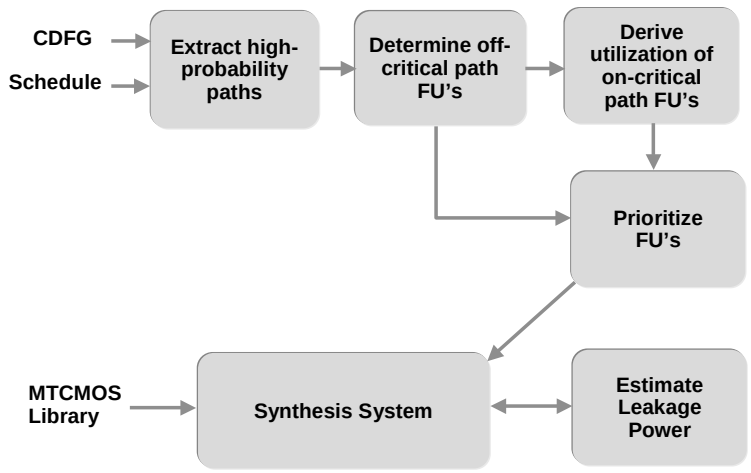

Leakage Results

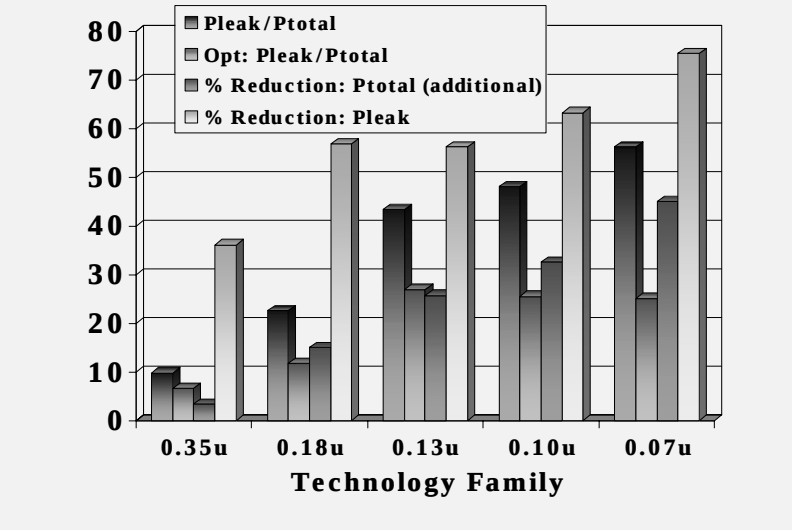

Hardware/software Co-Synthesis with DR-FPGAs

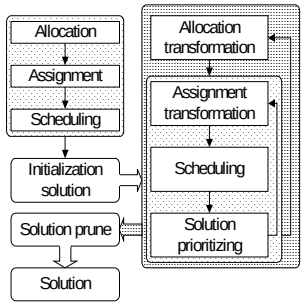

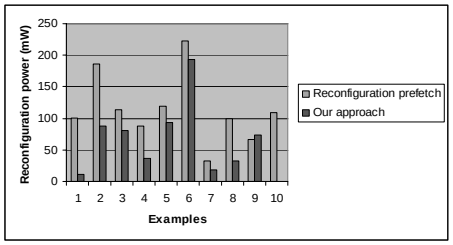

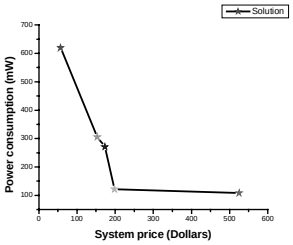

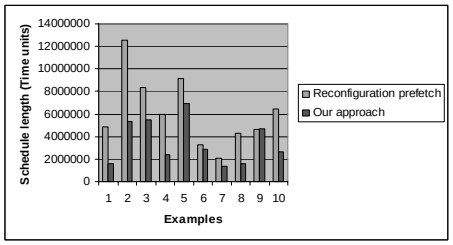

M. Pedram and N. Jha

Apollo

System-level Functional Partitioning

Functional Partitioning Integrated with System Synthesis
- Homogeneous model with partitioner/synthesizer
- Integrated evolutionary algorithm
- Optimize power/price under area and real-time constraints

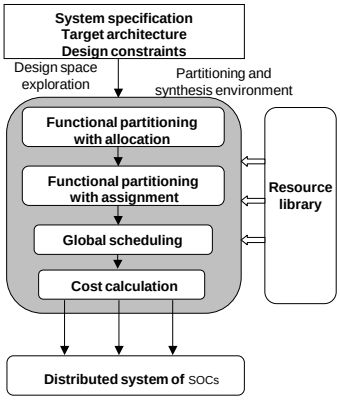

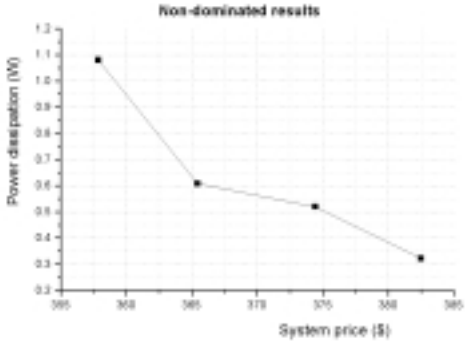

M. Pedram and N. Jha

Apollo

Summary II

- Input space adaptive SW synthesis reduces energy*delay by up to 66X
- Power-aware dynamic scheduler reduces power by up to 3.1X
- Battery-aware scheduler extends battery lifespan by up to 1.6X
- High-level SW energy macromodels speed up power estimation by 1-to-5 orders of magnitude with an average error of only 5%
- Leakage power optimization reduces leakage in .07 micron technology by up to 8.4X and total power by up to 2.3X
- HW-SW co-synthesis with DR-FPGAs reduces power by up to 6X
- Functional partitioning for distributed system of SOCs reduces power by up to 2.5X