Operating System and Algorithmic Techniques for Energy Scalable Wireless Sensor Networks

Amit Sinha and Anantha Chandrakasan

Massachusetts Institute of Technology

Overview

- The MIT µAMPS Project
  - Sensor Network and Node Model

Sensor Network

- Power-Aware Node Model
- Smart Shutdown
  - Hierarchical Sleep States

OS

- Energy-Quality Scalability
- Algorithmic Transformations
  - Filtering, Image Decoding Examples

Applications

The MIT µAMPS Project

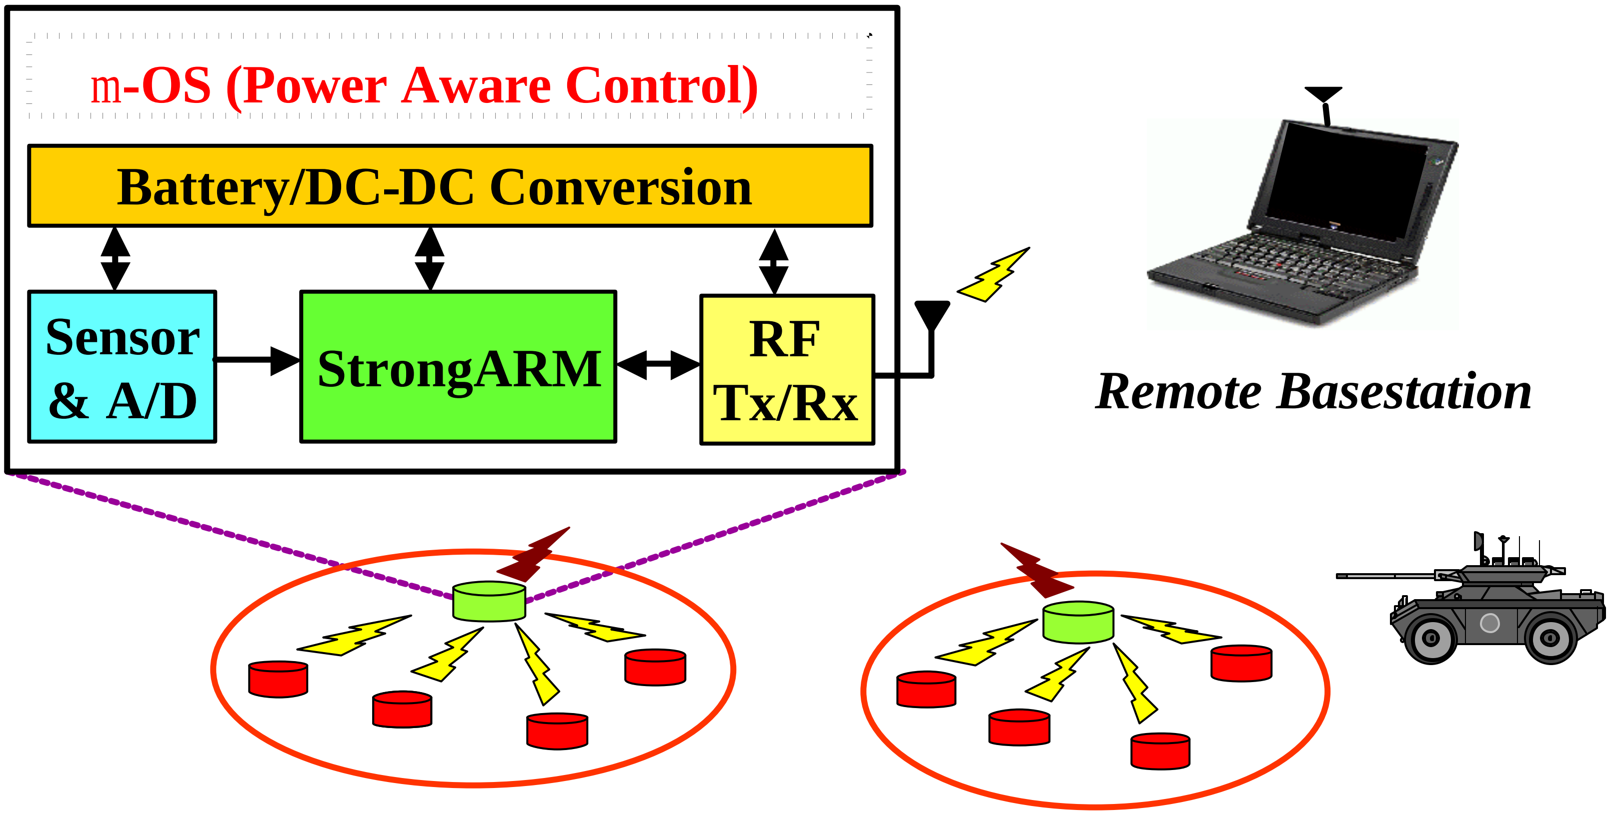

- A universal substrate for power aware data gathering from a massively distributed wireless network

Power Aware Node

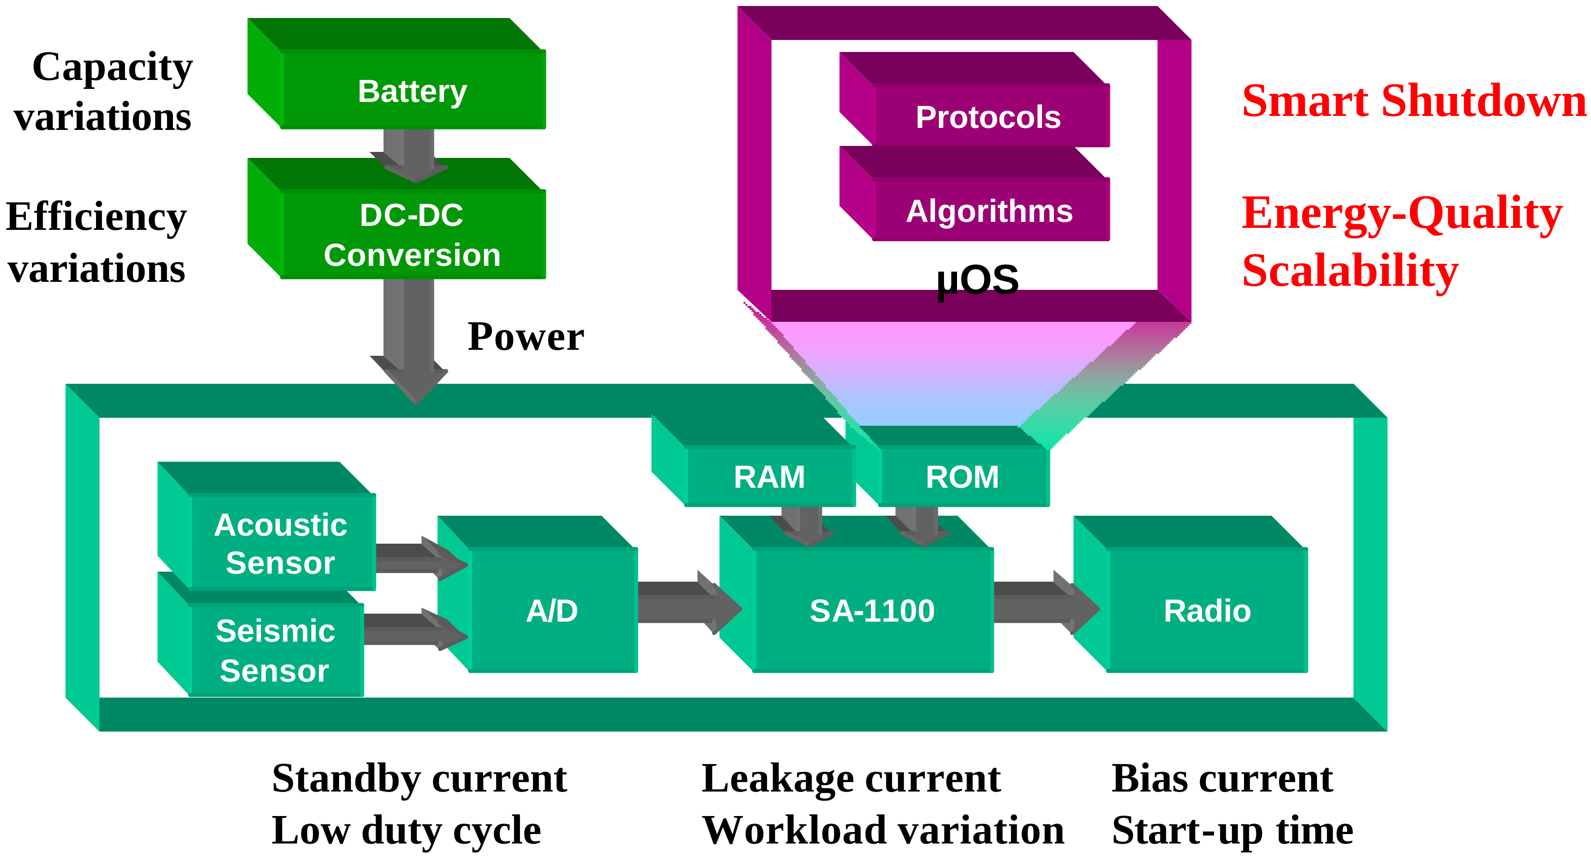

- Scalability ⇒ More efficient use of battery life

OS Directed Power Management

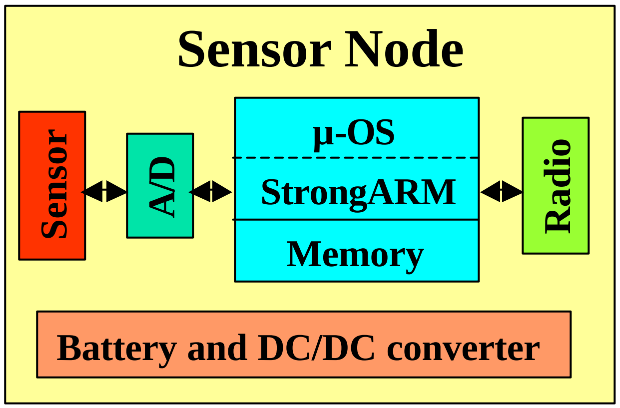

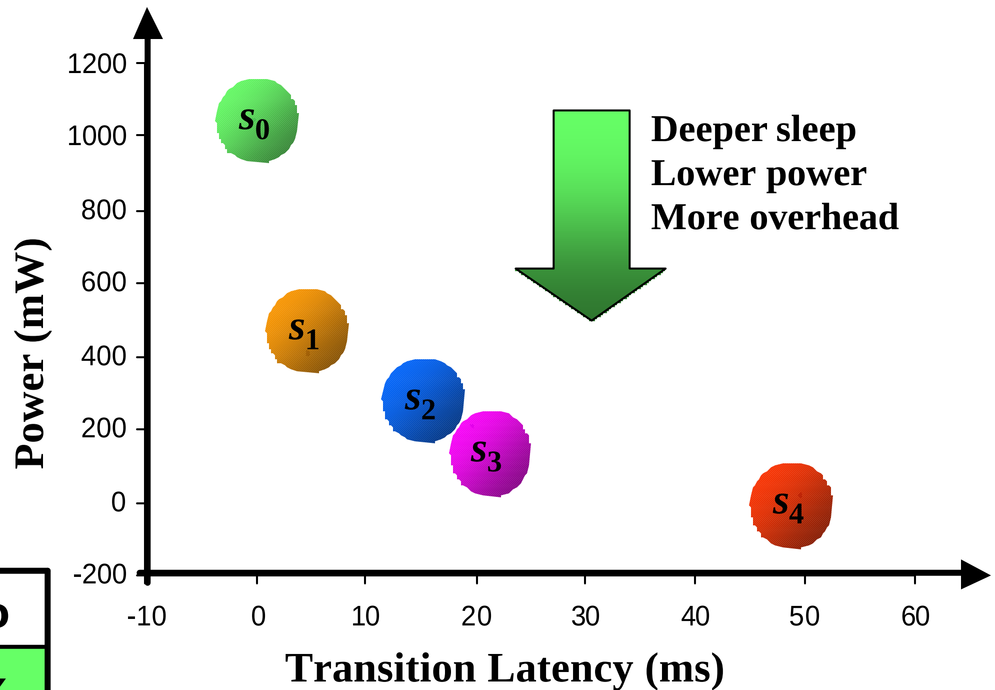

	ARM	Memory	Sensor	Radio
s_0	active	active	on	tx, rx
s_1	idle	sleep	on	rx
s_2	sleep	sleep	on	rx
s_3	sleep	sleep	on	off
s_4	sleep	sleep	off	off

- OS must decide suitable transition policy based on observed history

Event Model

- Poisson process
  - Rate : $?_{tot}$ s^{-1}
- Arbitrary spatial distribution
  - Probability of occurred event being at node$_k$ = p_{ek}
- Probability of no events at node$_k$ in time T_{th}

$$p_k(T_{th},0) = \sum_{i=0}^{\infty} e^{l_{tot}T_{th}} (l_{tot}T_{th})^i (1-p_{ek})^i = e^{-p_{ek} l_{tot} T_{th}}$$

- Probability of at least 1 event

$$p_{th,k}(T_{th}) = 1 - e^{-p_{ek} l_{tot} T_{th}}$$

Sleep State Transitions

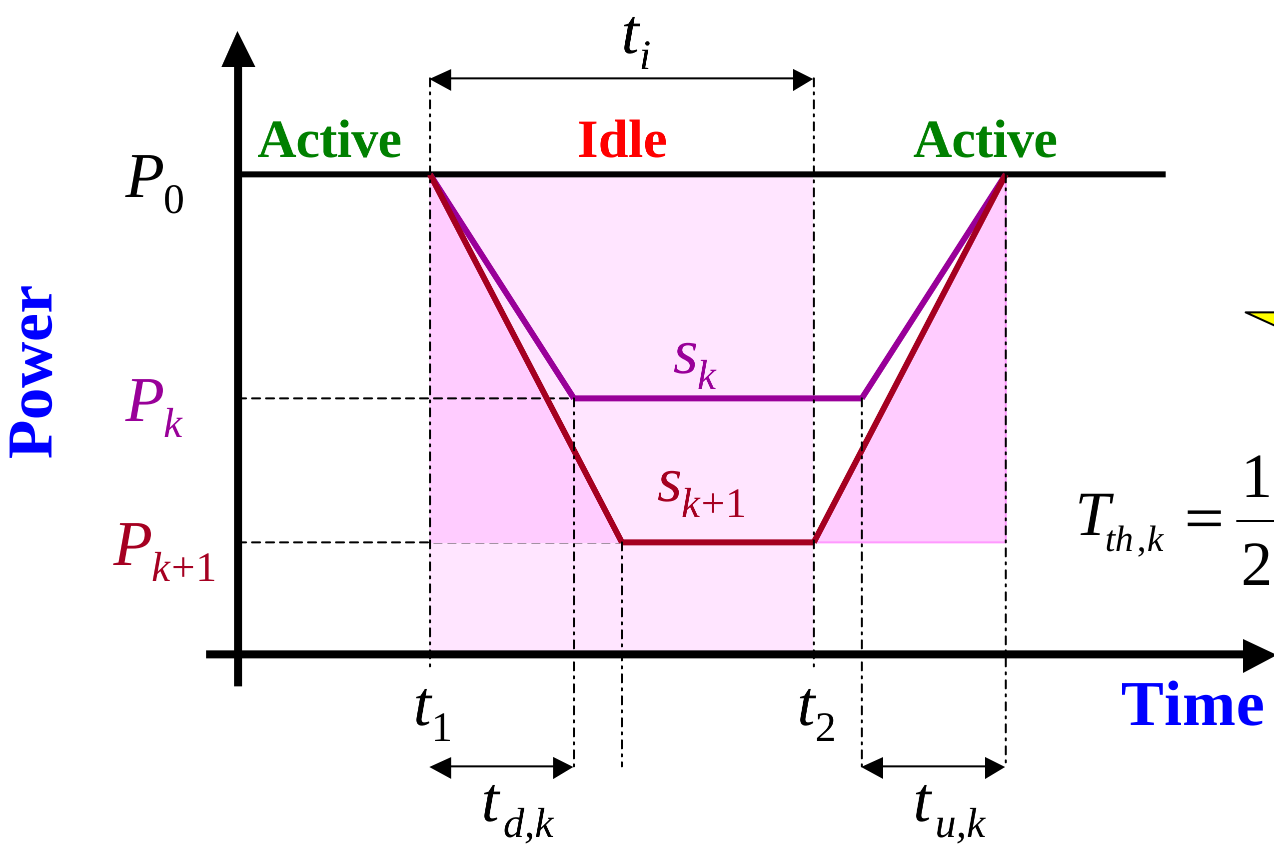

More idle time before energy savings is positive

$$T_{th,k} = \frac{1}{2}\left[t_{d,k} + \left(\frac{P_0 + P_k}{P_0 - P_k} \right) t_{u,k} \right]$$

- Every state has a minimum time t_i before which energy savings is negative!

Typical Thresholds and Policy

	Power (mW)	Transition Time (ms)	$T_{th,k}$ (ms)
s_0	1040	-	-
s_1	400	5	8
s_2	270	15	20
s_3	200	20	25
s_4	10	50	50

- Transition to s_k if probability of at least one event is greater than a constant (p_{th0})

Missed Events

- State s_4 is 'completely off'
  - Disallow s_4 – Critical sensing
  - Selectively disallow s_4
    - If idle > threshold, sleep with probability p_{s4}
    - Wake up after a time

$$t_{s4,k} = -\frac{1}{l_k}\ln(p_{s4})$$

- Overall 2 system parameters
  - p_{th0} and p_{s4}

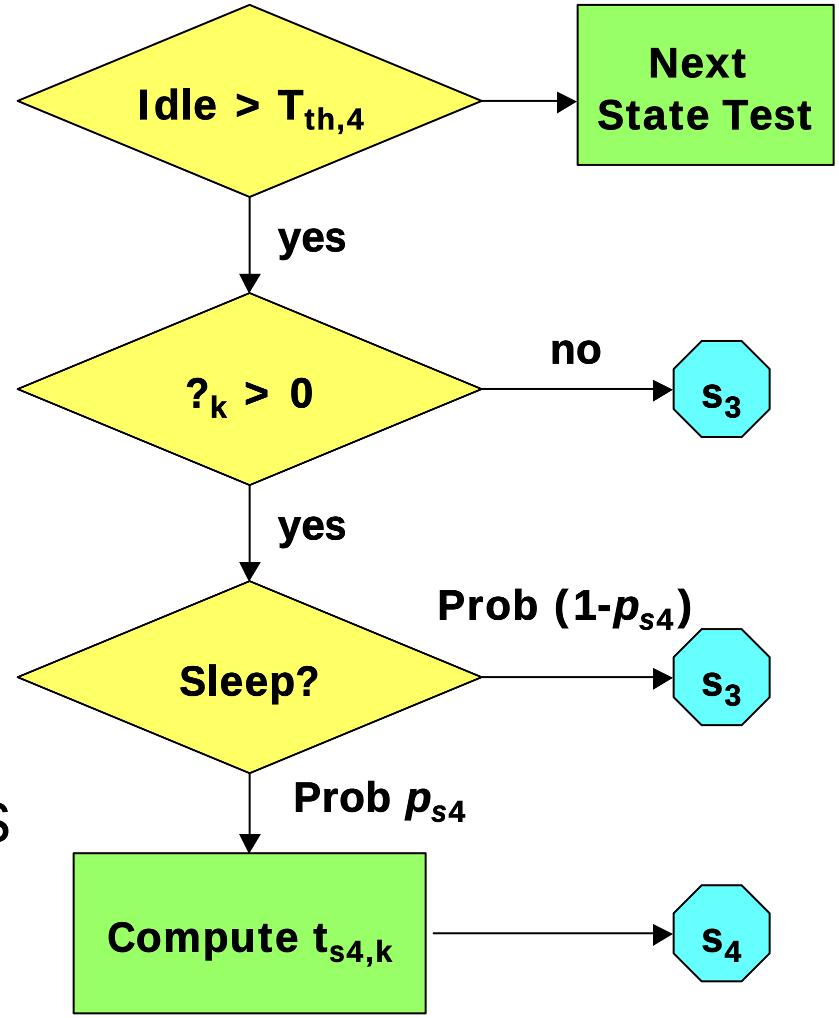

Simulation Infrastructure

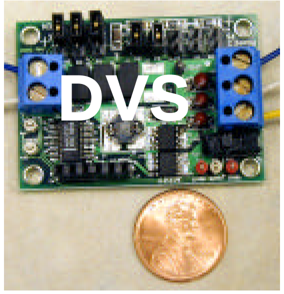

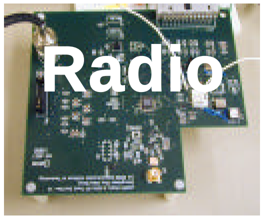

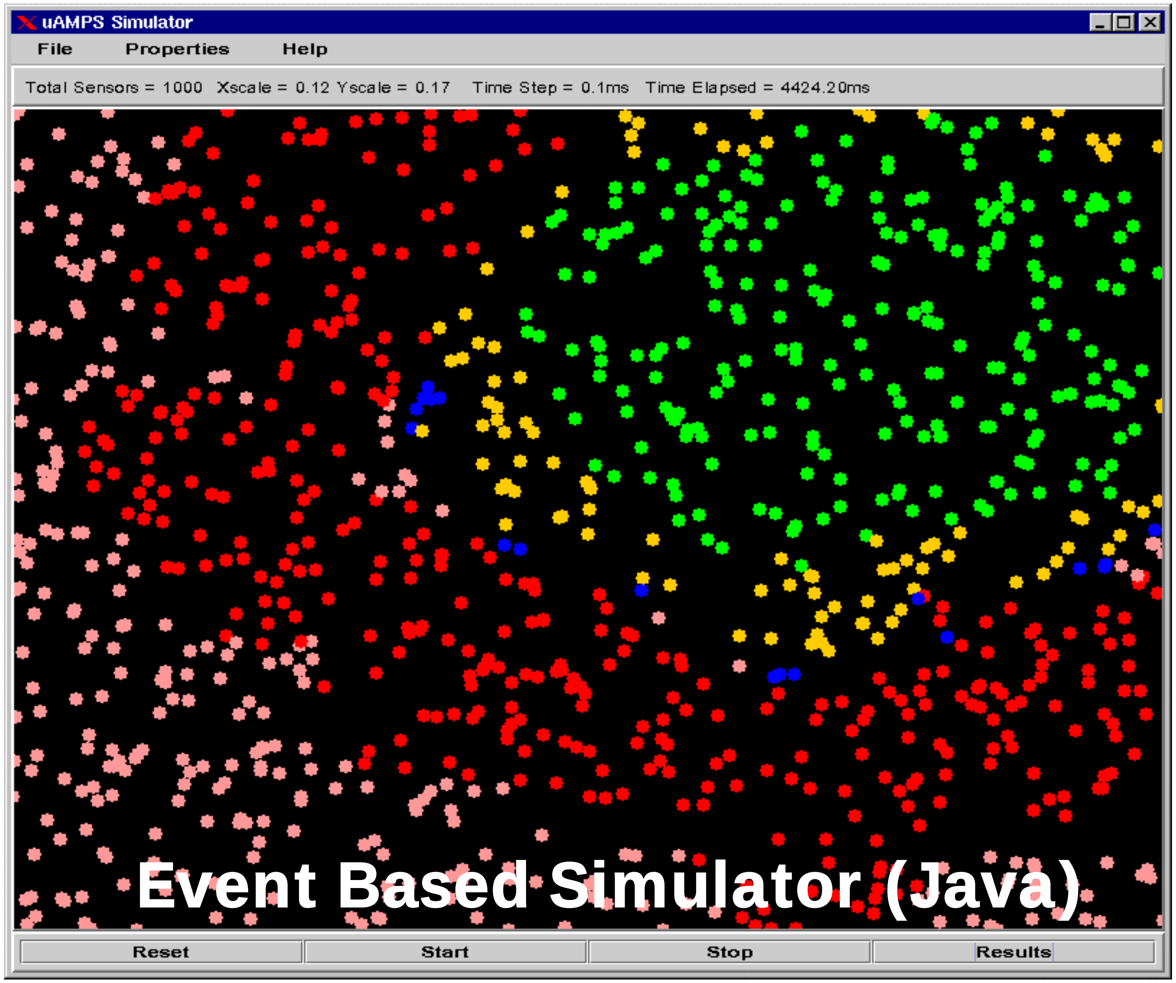

- Node models from actual hardware

Results

Spatial Event Distribution

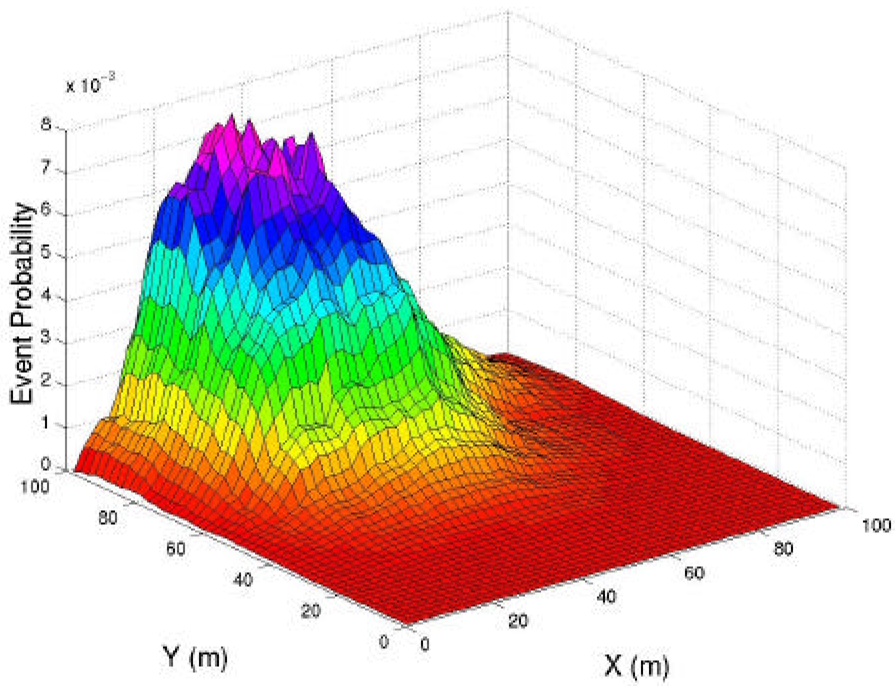

Spatial Energy Consumption

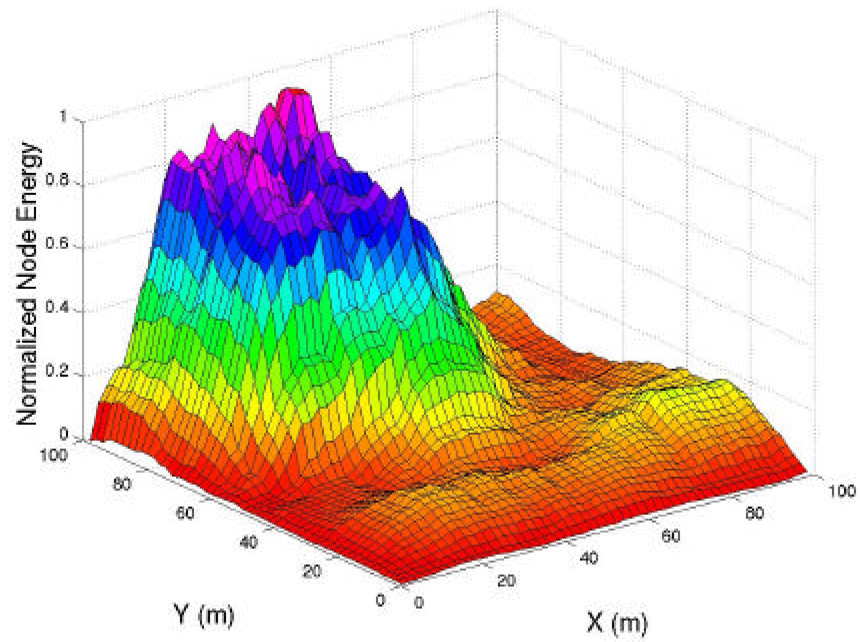

- Energy Consumption Tracks Events

Results

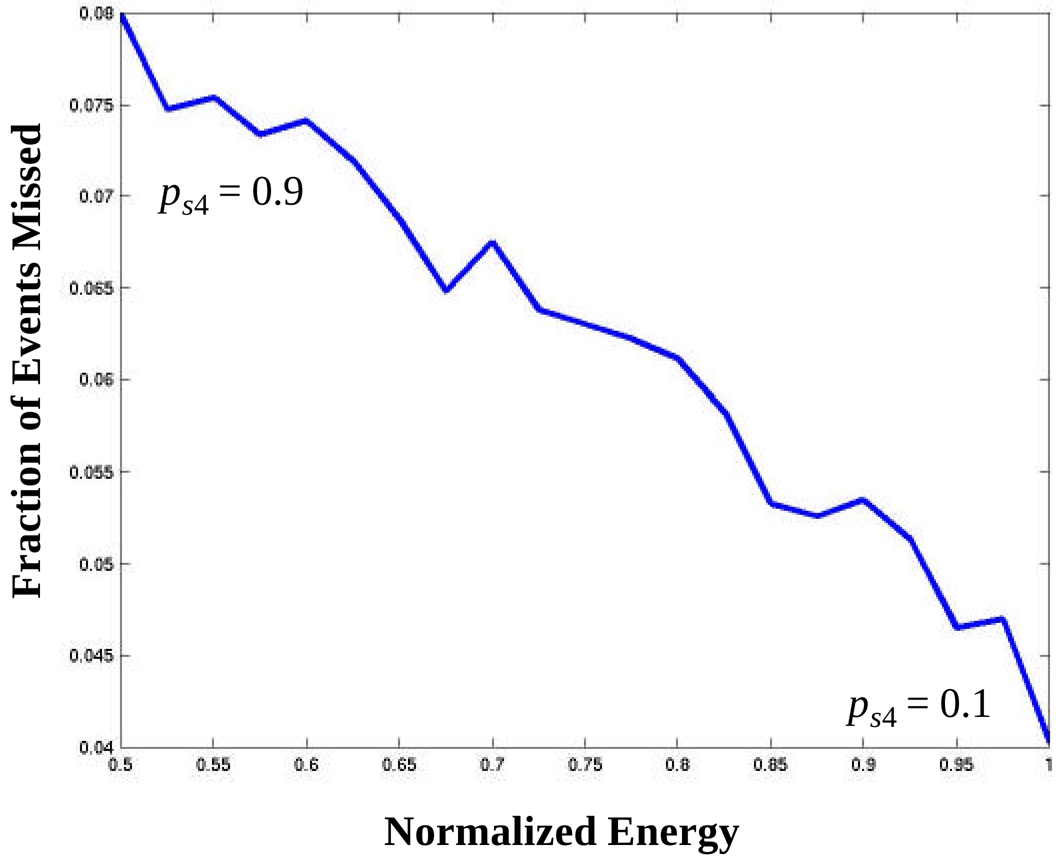

- Scalable degradation in event detection

Energy Scalable Applications

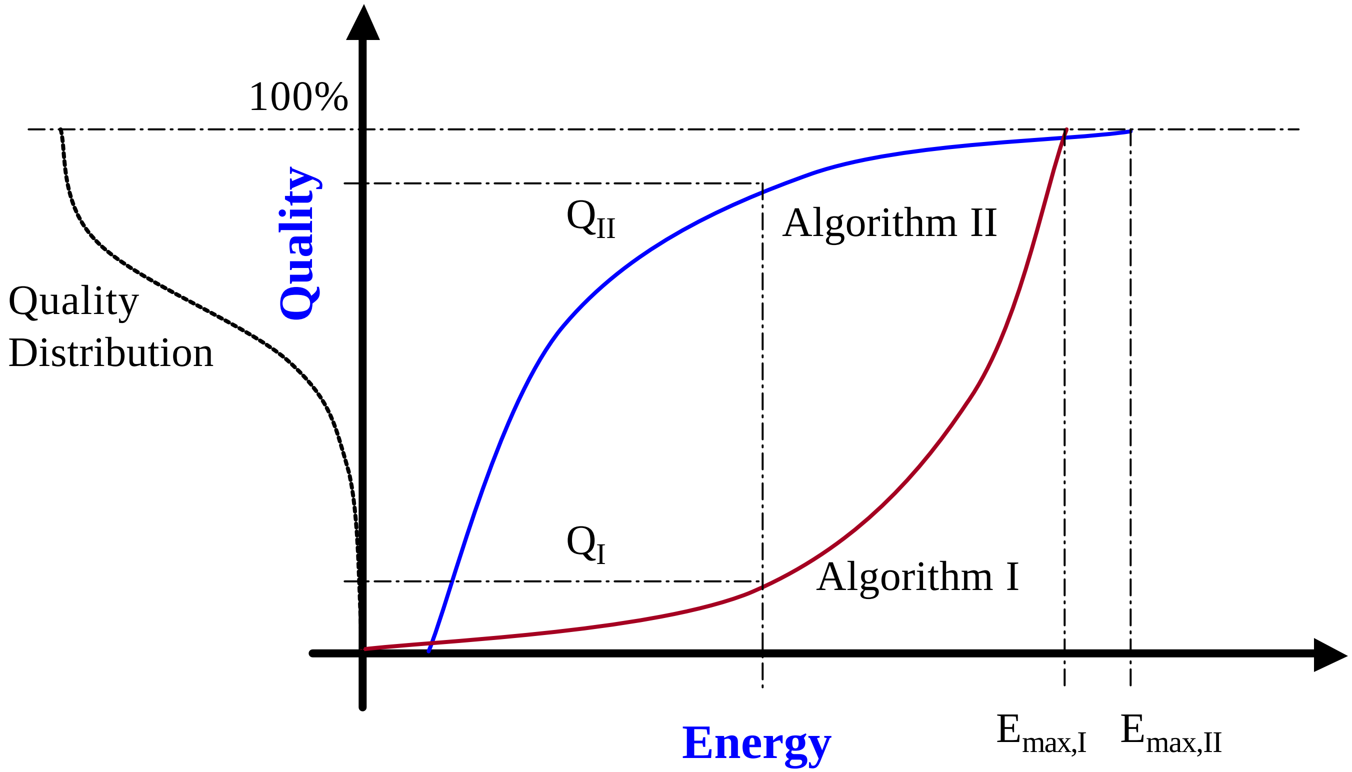

- Maximize quality for given energy availability
- Energy Quality (E-Q) graph maximally concave

Simple Example

$$y = f(x) = 1 + k_1 x + k_2 x^2 + \cdots + k_N x^N$$

Original	Scalable

`xpowi = 1.0;`
`y = 1.0;`
`for(i=1; i<N; i++) {`
` xpowi *= x;`
` y += xpowi*k[i];`
`}`
	`if(x > 1.0) {`
	` xpowi = pow(x,N);`
	` y = k[N]*xpowi+1;`
	` for(i=N-1; i>0; i--) {`
	`  xpowi /= x;`
	`  y += xpowi*k[i]; }`
	`}`
	`else { //original algo }`

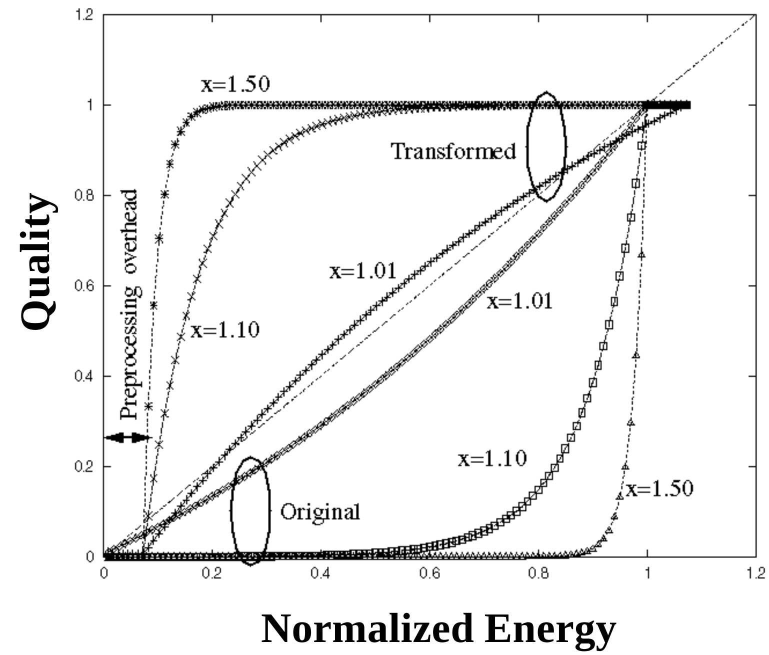

- Incremental refinement
- Most-significant-first approach

Scalable Filtering

$$y[n] = \sum_{k=0}^{N-1} h[k]x[n-k]$$

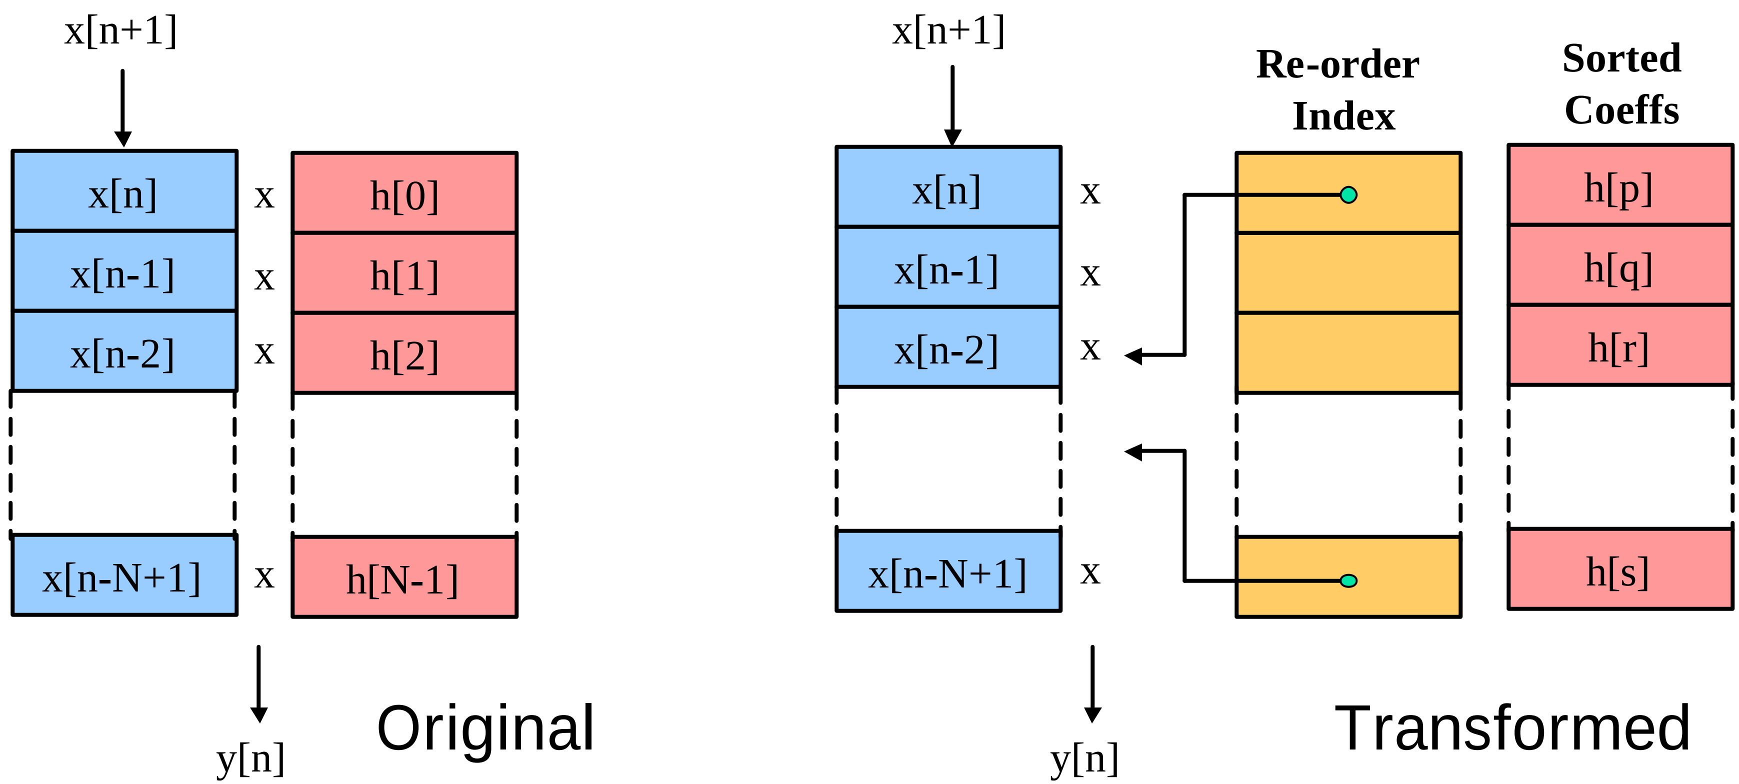

Scalable Filtering E-Q

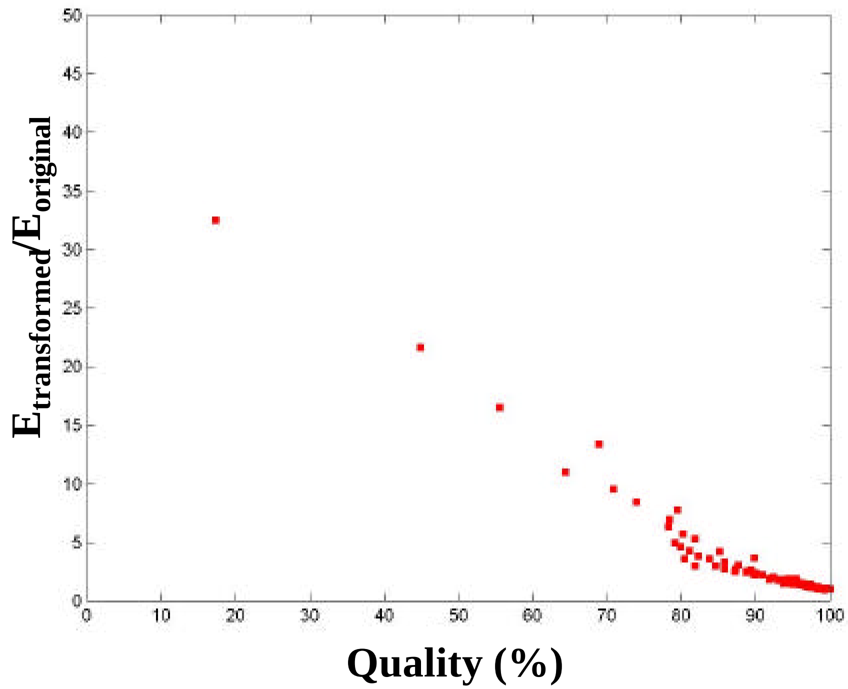

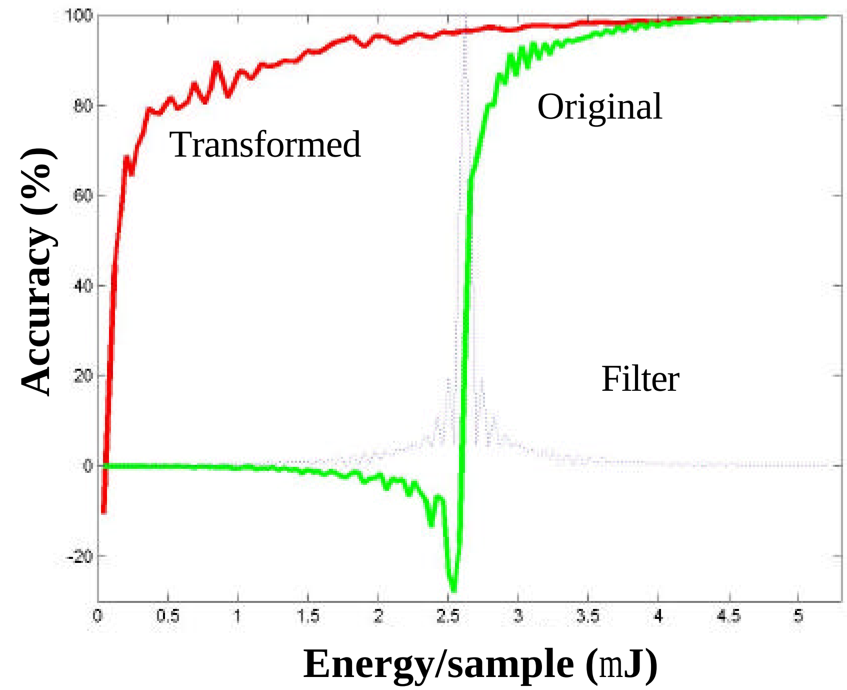

- 128 Tap FIR filtering on speech data
- StrongARM measurements
  - 5.12µJ per sample
  - 0.21µJ per sample overhead (4%)

Scalable Image Decoding: IDCT

$$x[i,j]=\frac{1}{4}\sum_{u=0}^{7}\sum_{v=0}^{7}c[u]c[v]X[u,v]\cos\left(\frac{(2i+1)u\pi}{16}\right)\cos\left(\frac{(2j+1)v\pi}{16}\right)$$

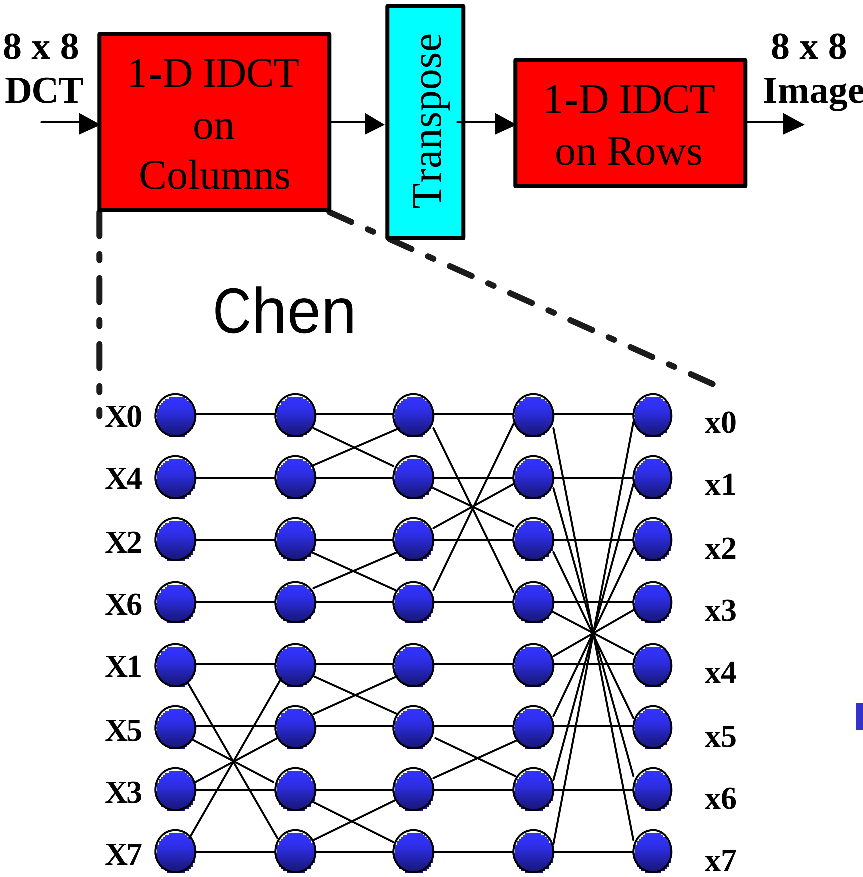

$$\begin{bmatrix} x_{0,0} \\ x_{0,1} \\ x_{0,2} \\ \vdots \\ \vdots \\ x_{7,7} \end{bmatrix} = X_{0,0}\begin{bmatrix} c_0^{0,0} \\ c_1^{0,0} \\ c_2^{0,0} \\ \vdots \\ \vdots \\ c_{63}^{0,0} \end{bmatrix} + X_{0,1}\begin{bmatrix} c_0^{0,1} \\ c_1^{0,1} \\ c_2^{0,1} \\ \vdots \\ \vdots \\ c_{63}^{0,1} \end{bmatrix} + \cdots + X_{7,7}\begin{bmatrix} c_0^{7,7} \\ c_1^{7,7} \\ c_2^{7,7} \\ \vdots \\ \vdots \\ c_{63}^{7,7} \end{bmatrix}$$

FM-IDCT

- Scalability at the cost of slightly more operations

Most Significant DCT Coefficients

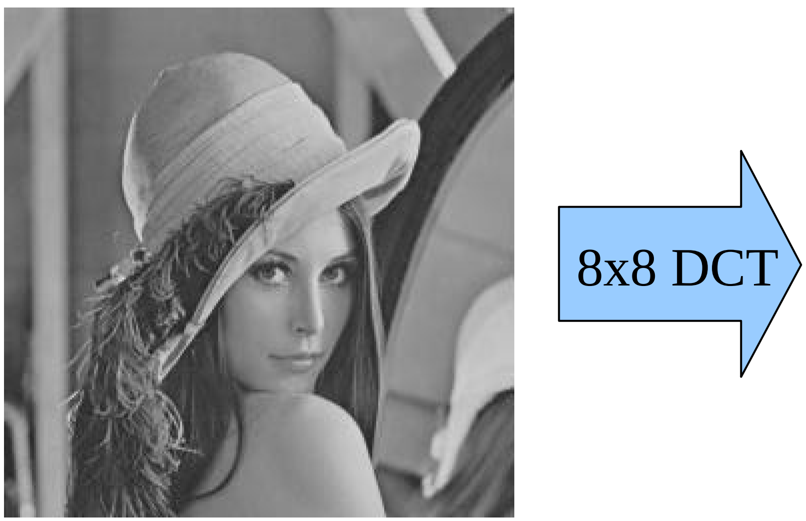

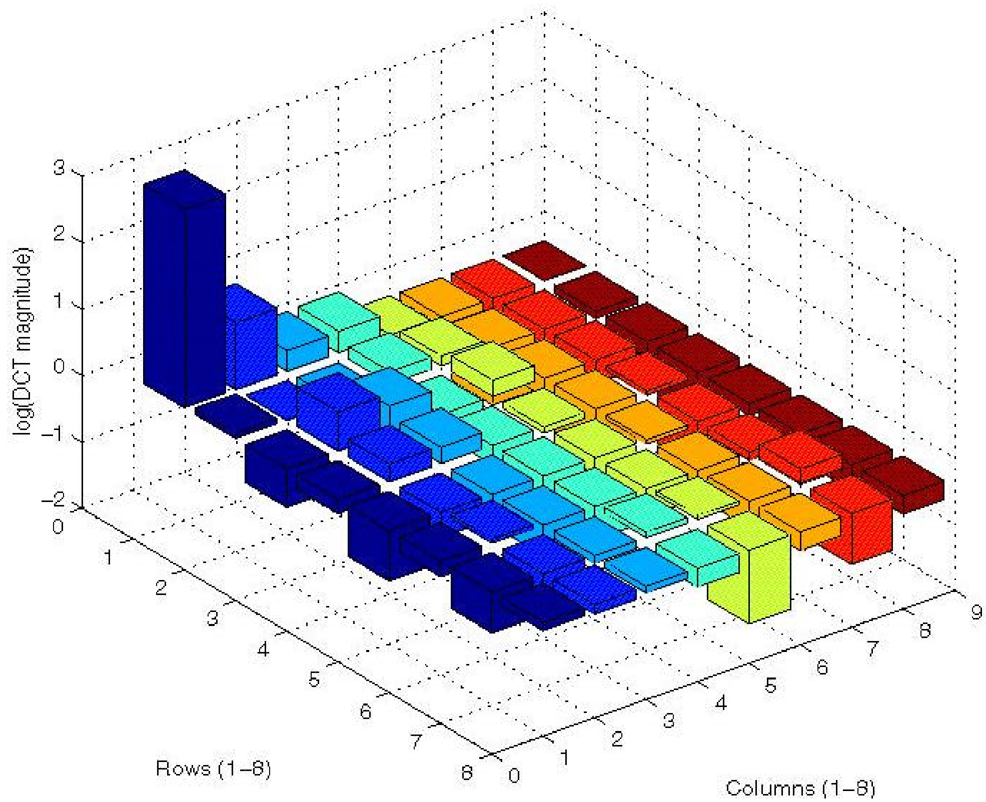

- Most energy concentrated in lower coefficients
- Accumulate lower frequency components first

Scalable vs Non-Scalable IDCT

Chen

FM-IDCT

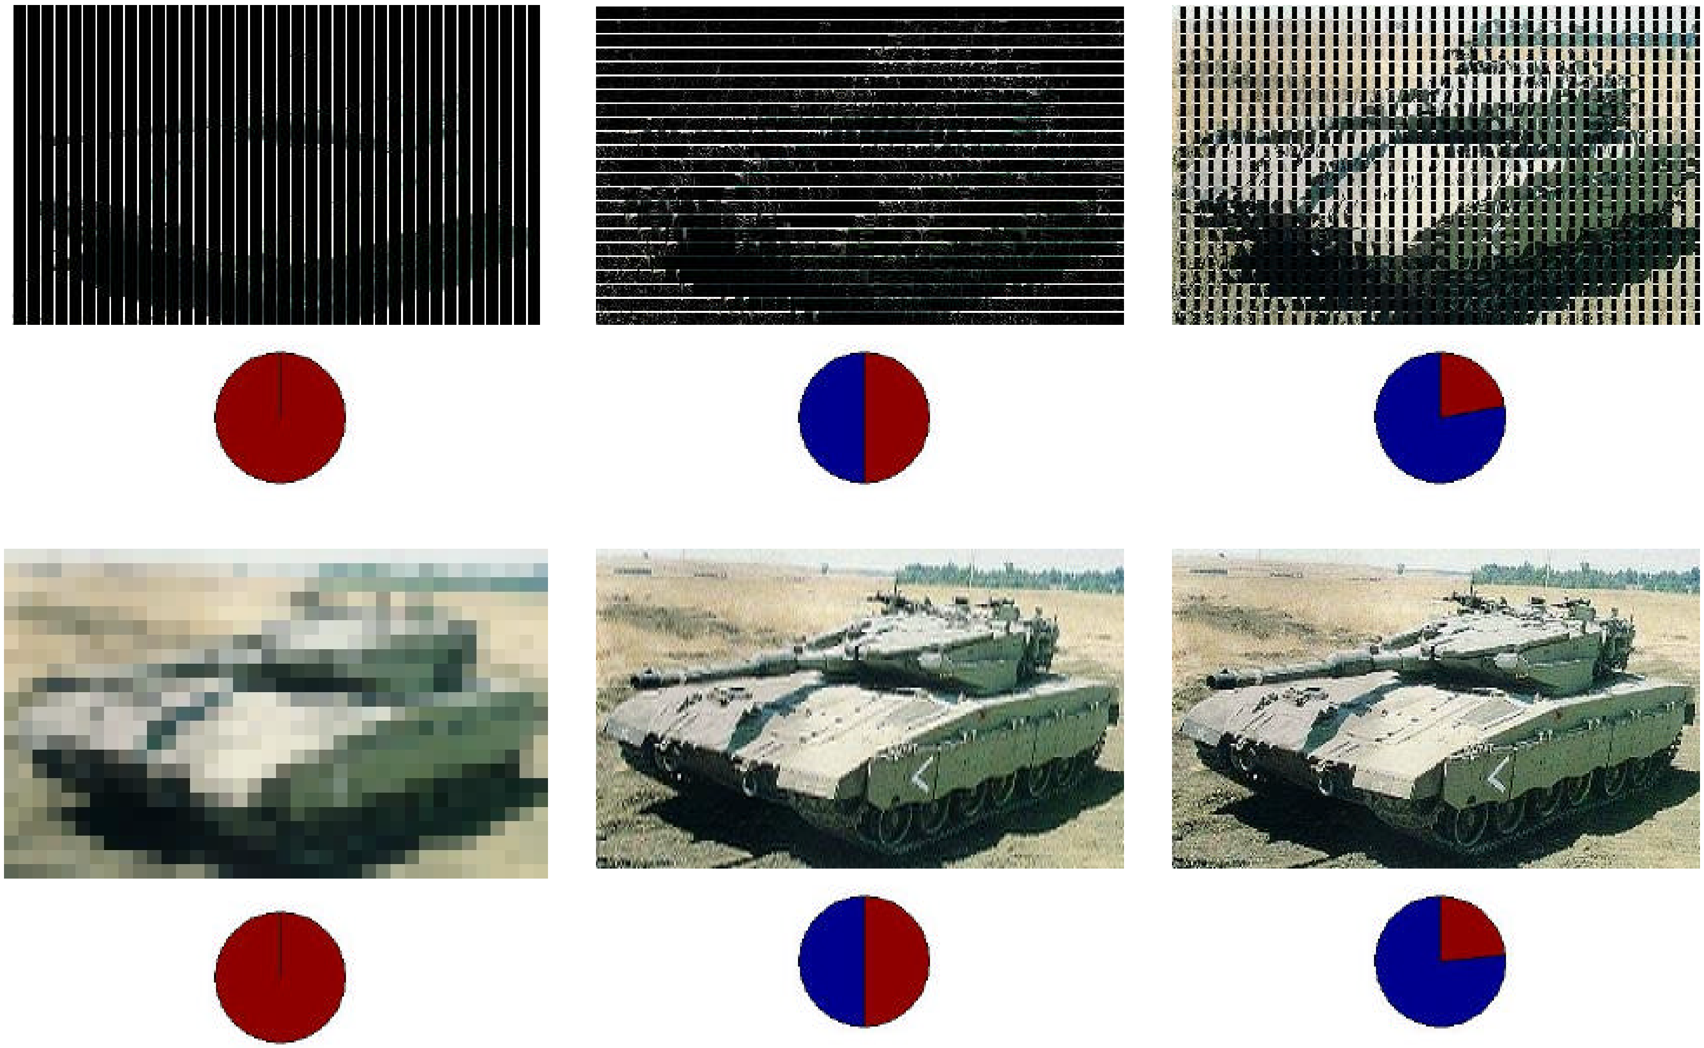

IDCT E-Q

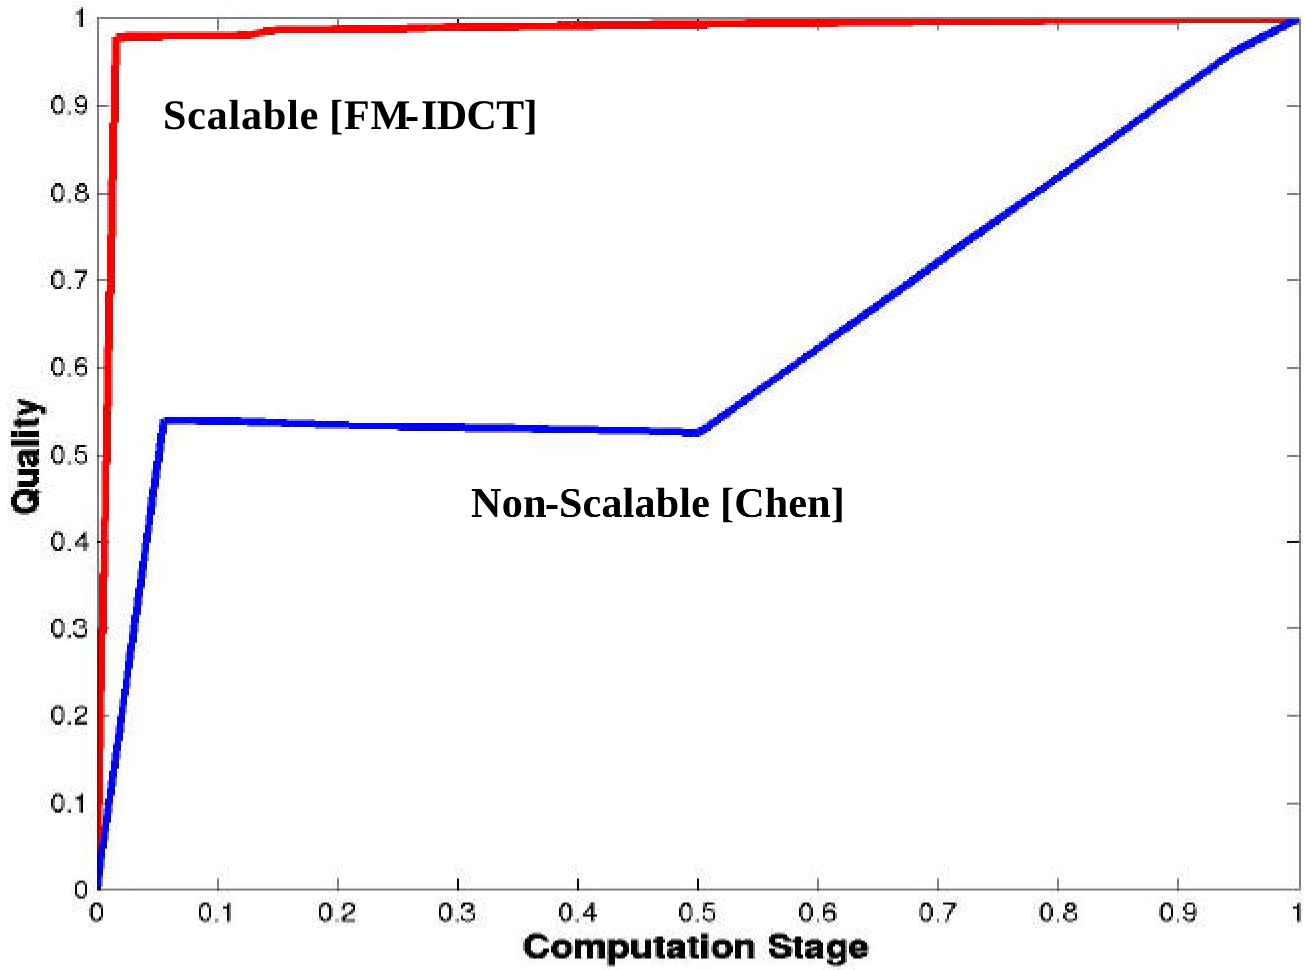

Scalable [FM-IDCT]
Non-Scalable [Chen]
Quality
Computation Stage
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1

Conclusions

- Power Management in Sensor Networks using
  - OS directed smart hierarchical shutdown
  - Scalable Algorithms
- Power Aware Sensor Node and Network
  - Power and transition overheads
  - Energy gain idle thresholds
- Energy-Quality Scalability
  - Missed events vs. Energy
  - Algorithmic transformations for scalability