Power-Aware Systems

Manish Bhardwaj, Rex Min and Anantha Chandrakasan

Massachusetts Institute of Technology

November 2000

Power-awareness: Intuitive Notions

- **Motivation: Maximize lifetime of energy constrained systems**
  - » **Maximize system-level energy efficiency**
- **Implication: Given an operating scenario, consume only as much energy as the scenario demands**
- **Alternately, scale the power consumed in response to changing scenarios (power-awareness)**

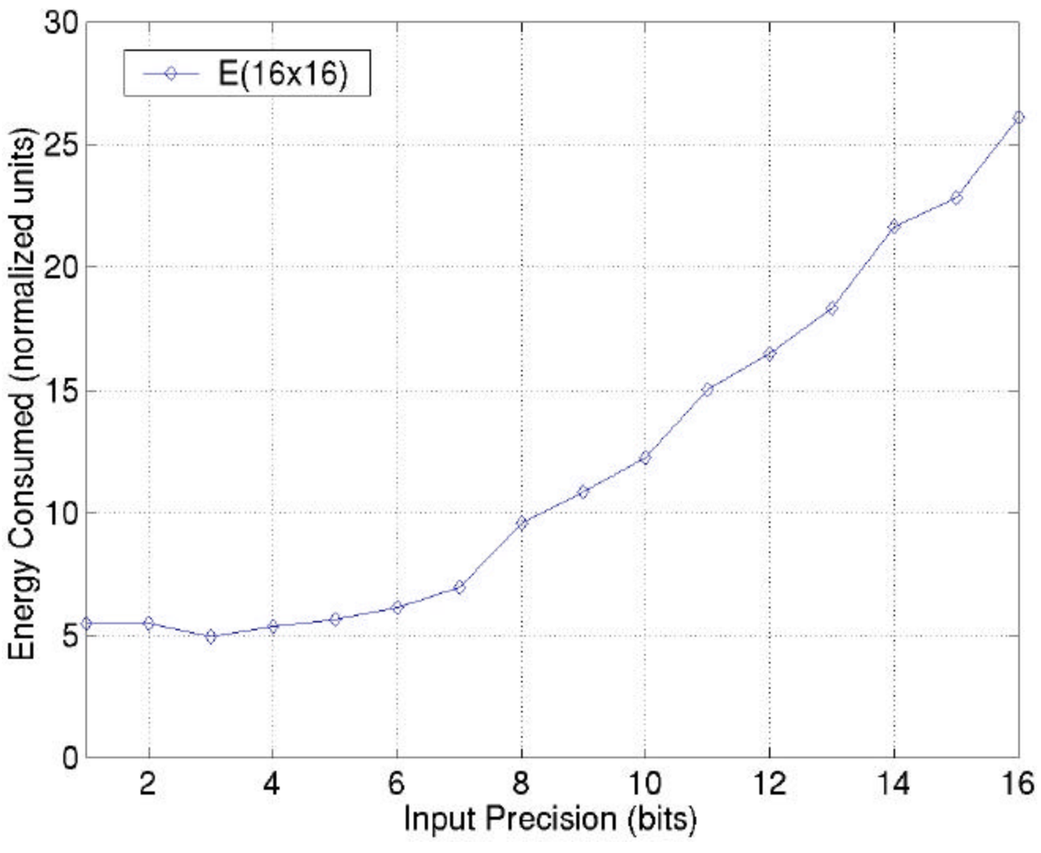

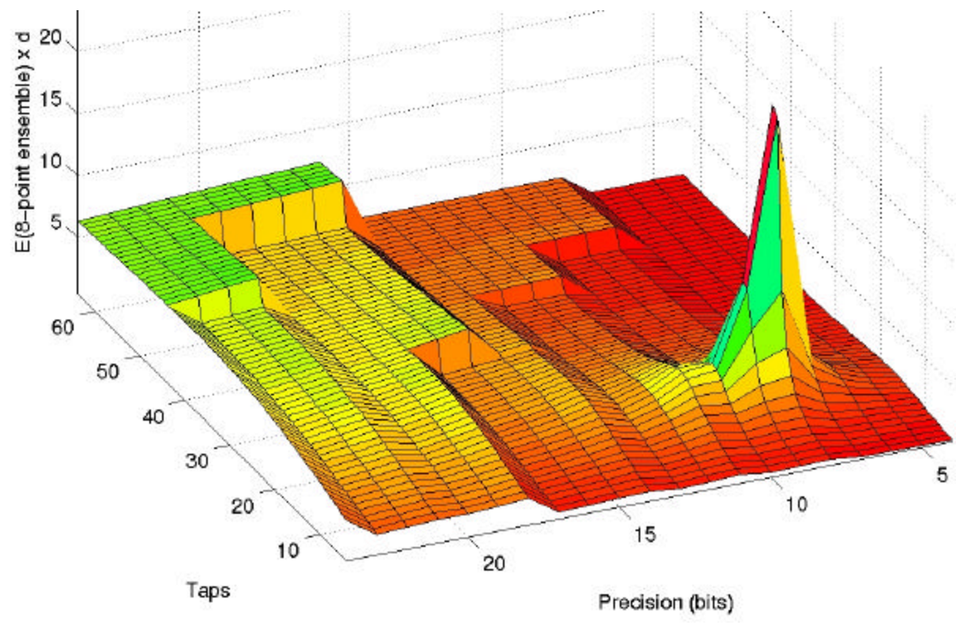

Agenda

- **Key questions**
  - **What are operating scenarios?**
  - **How *well* are these systems tracking their scenarios?**
  - **What can we do to *improve* this tracking?**
  - **What are the *costs* and *benefits*?**
- **Abstractions**
  - **Awareness dimensions, operating scenarios, energy curves, scenario distributions**
- **Formalizing Power-Awareness**
- **Enhancing Power-Awareness**
- **Examples:**
  - **Multipliers**
  - **Register Files**
  - **Filters**
  - **Analog-Digital Converters**
  - **Variable-Voltage Processors**
  - **Wireless Networks**

Abstractions: Scenarios

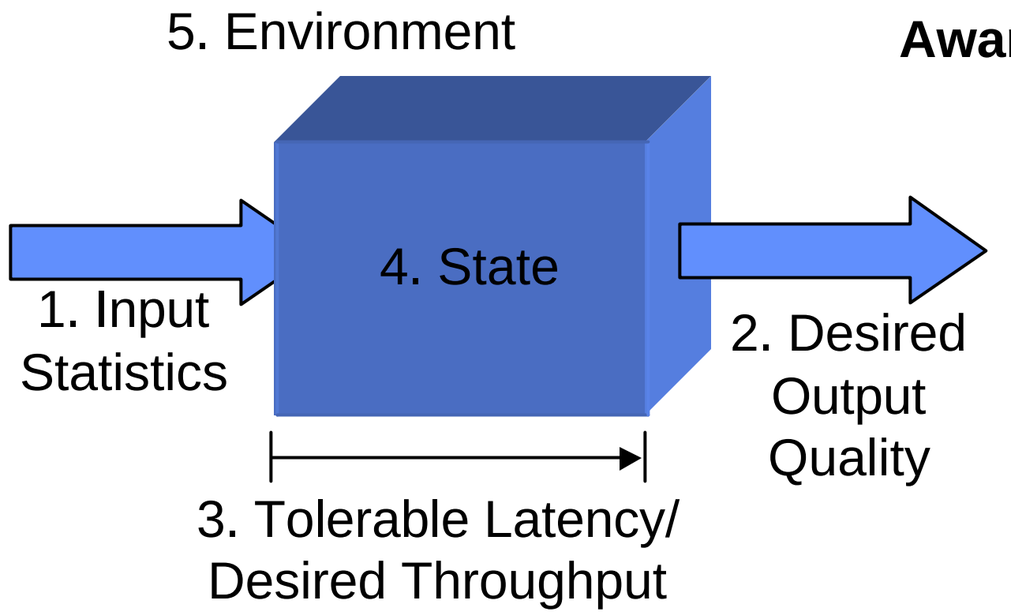

- **Over any specified time interval, the energy consumed by a system is governed by five key dimensions**
- **Scenarios are characterized by precisely these dimensions**
- **Scenario ° <Input, Output Quality, Latency, State, Environment>**
- **Choices in specifying scenarios**
  - **Number of dimensions to include**
  - **Detail with which the dimension is captured**
- **Example: Characterizing scenarios in a 16x16-bit multiplier**

Scenario Characterization in Multipliers

- **Input dimension only**
  - **Scalar *m*: Specifies a maximum precision requirement**
  - **Unordered pair (*m*, *n*): Specifies a *m*x*n*-bit multiplication**
  - **Ordered pair <*m*, *n*>**
  - **Ordered operands <X,Y>**
- **Input and state**
  - **Ordered operands and previous operands <X[n],Y[n],X[n-1],Y[n-1]>**
- **Input, state and desired precision**
- **Input, state, desired precision and latency**

Abstractions: Energy Curves

- **The energy consumed by a system as a function of its scenario, *E*(*H*, *s*)**

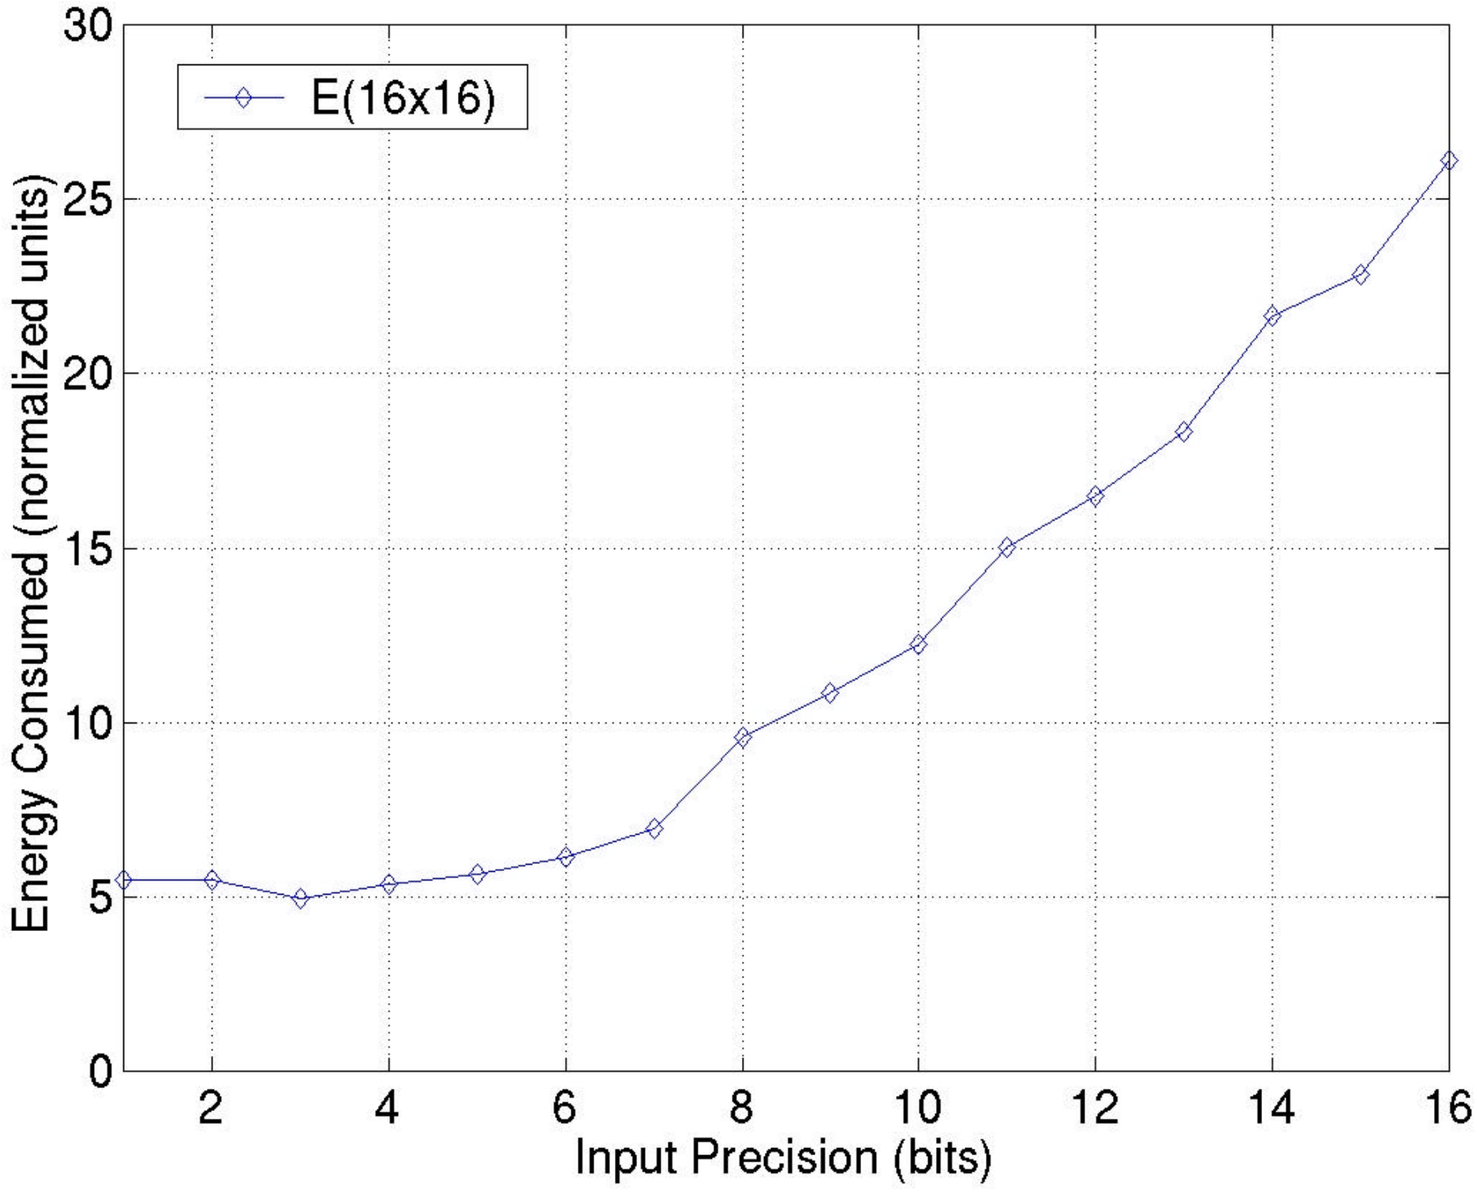

Abstractions: Scenario Distributions

- The probability that a system will reside in a certain scenario is captured by scenario distributions, $d_S(s)$

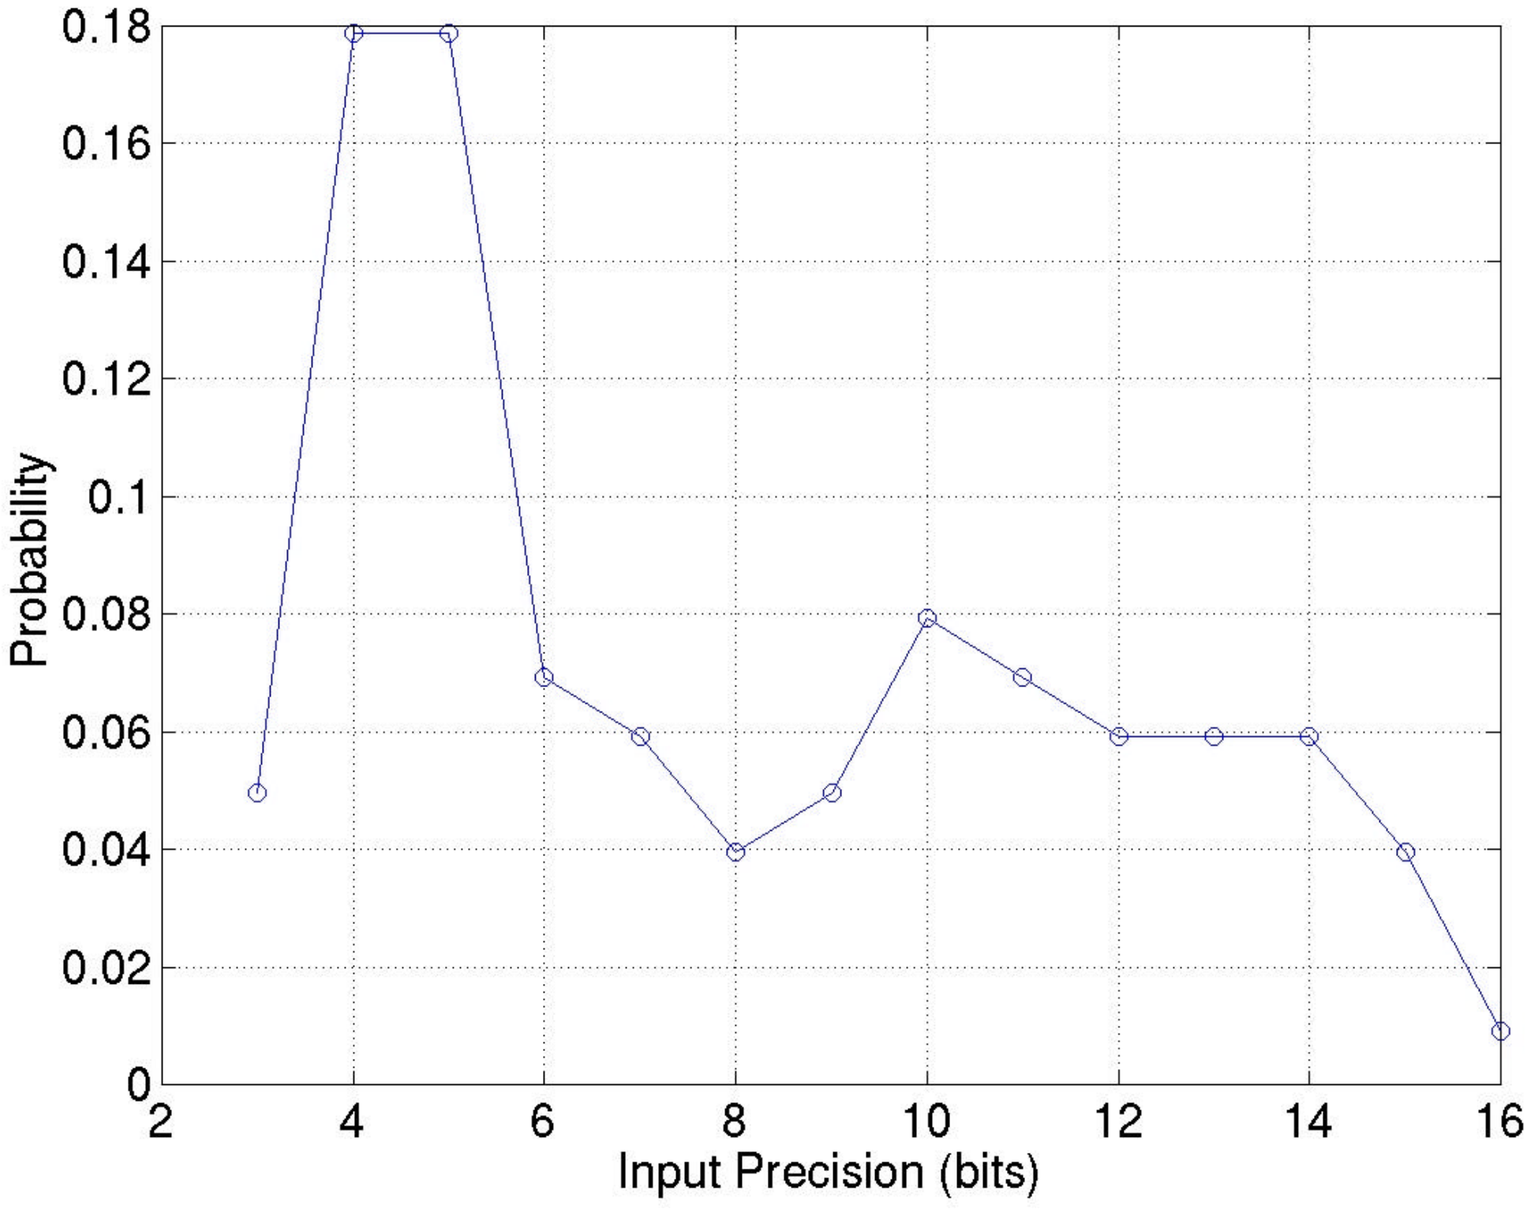

Perfect Power Awareness

A system is termed perfectly power-aware iff it consumes only as much energy as its current scenario demands.

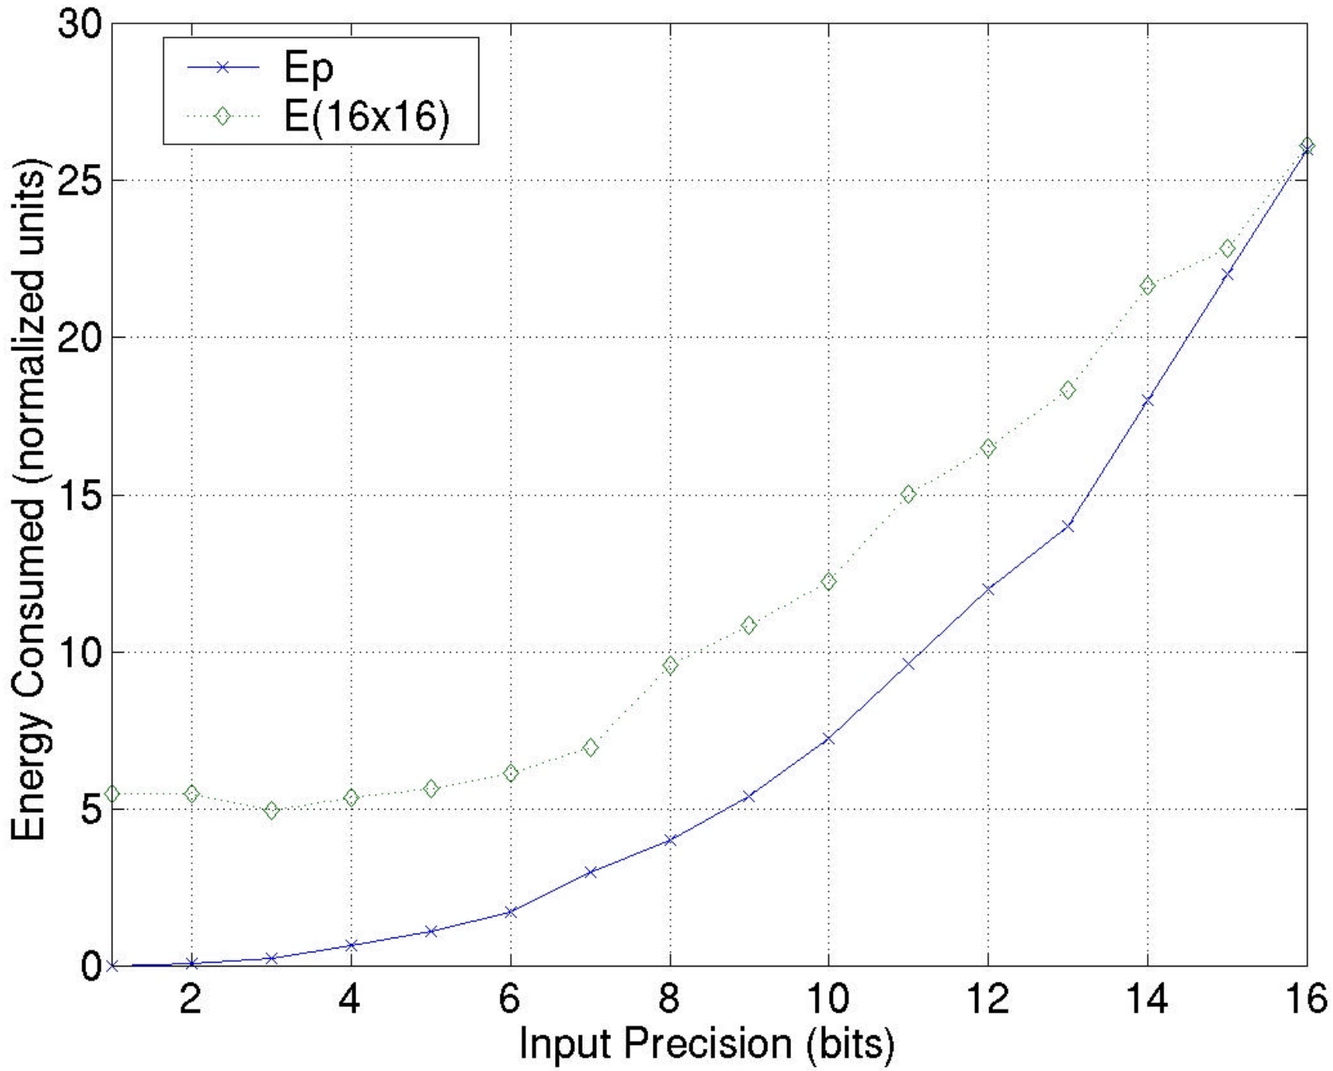

- **Perfect energy curve obtained by constructing dedicated point systems**

Perfect Systems

- **A system that would result in $E_{perfect}$ is termed the perfect system ($H_{perfect}$)**

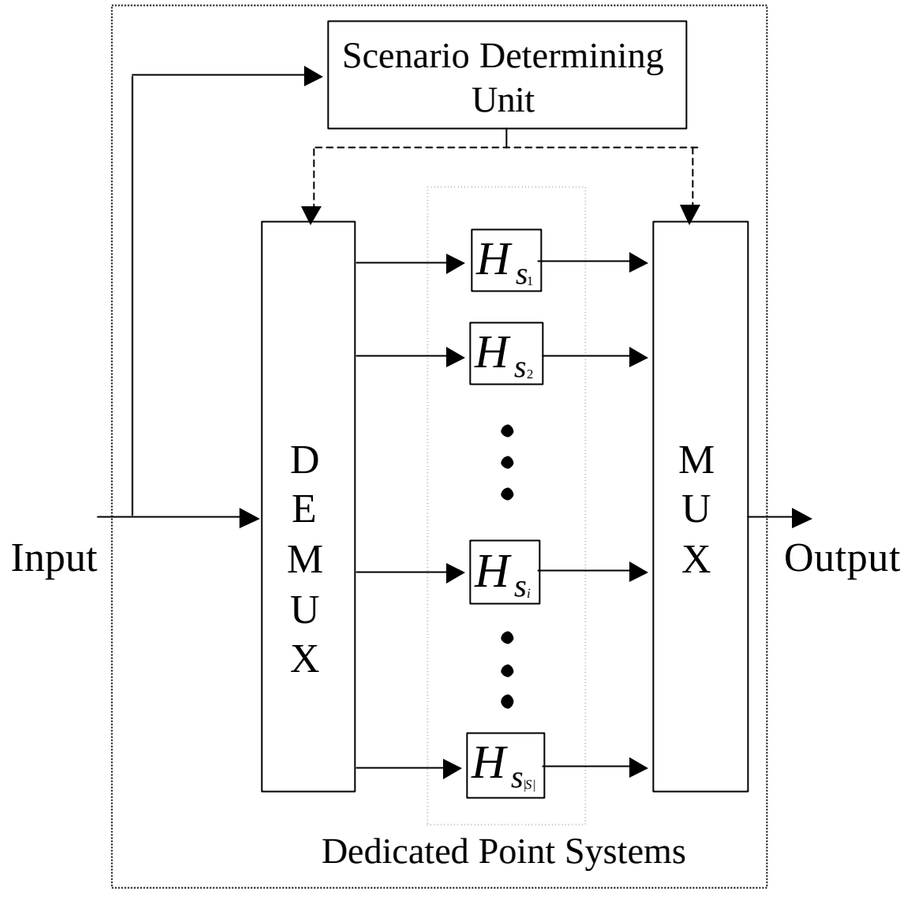

- **If scenario detection and interconnect costs were zero, the system above would yield $E_{perfect}$**

Quantifying Power Awareness

- **The relative energy curve is simply the energy curve of a system normalized to the perfect energy curve**

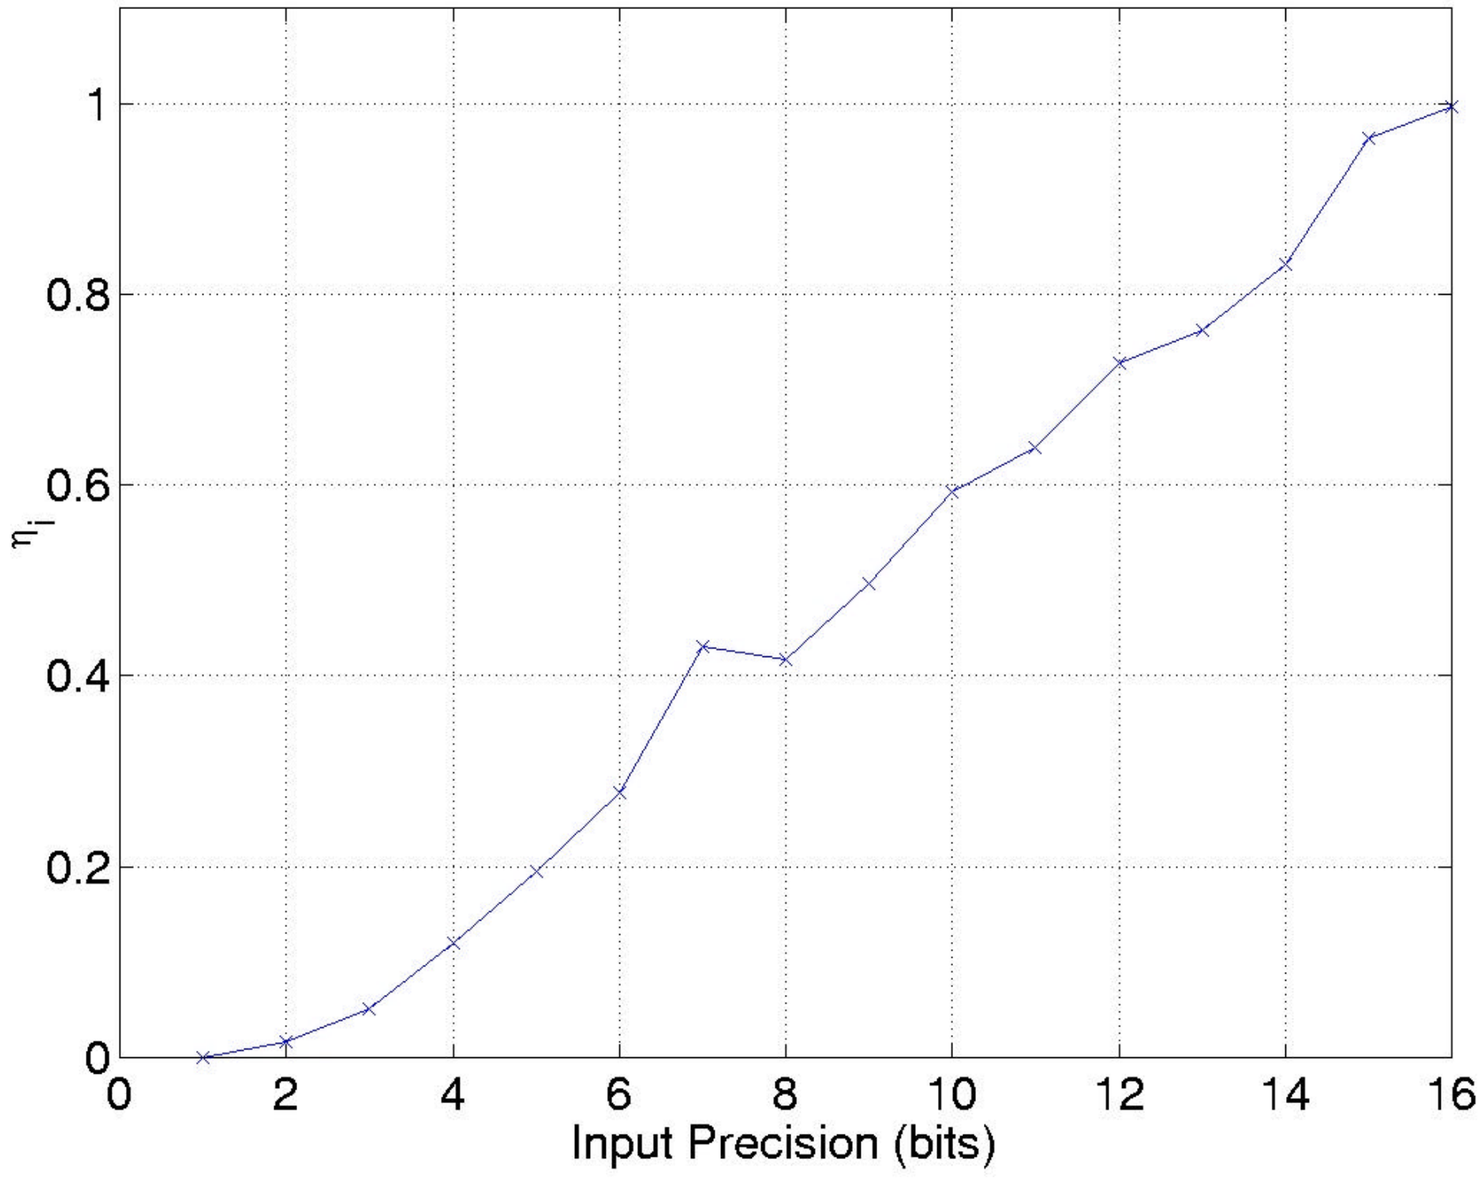

Power Awareness Metric

- **Reduce the relative curve to a single number by appropriate weighting**
  - **Weigh by probability of occurrence of scenario**
  - **Weigh by energy dissipated in the scenario**

$$f = \left[\frac{\sum_{i \in Scenarios} h(s_i) d_S(s_i) E(H, s_i)}{\sum_{i \in Scenarios} d_S(s_i) E(H, s_i)} \right] = \left[\frac{\sum_{i \in Scenarios} E(H, s_i) d_S(s_i)}{\sum_{i \in Scenarios} E(H_{perfect}, s_i) d_S(s_i)} \right]^{-1}$$

- **Physical interpretation: Expected system lifetime normalized to lifetime of perfect system**
- **Defined w.r.t scenario distribution and a set of point systems**
- **Metric leads to complete ordering for a specified distribution and partial ordering otherwise**

Enhancing Power-Awareness: Ensemble Construction

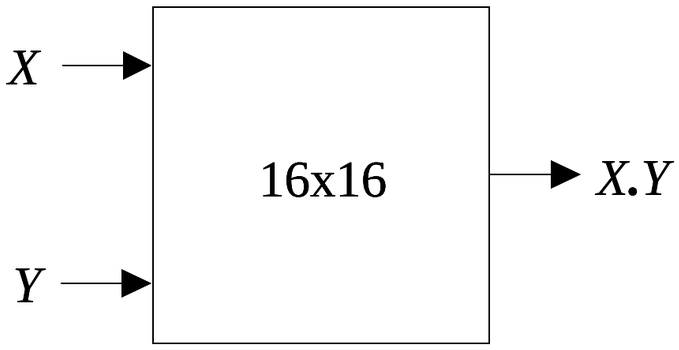

versus

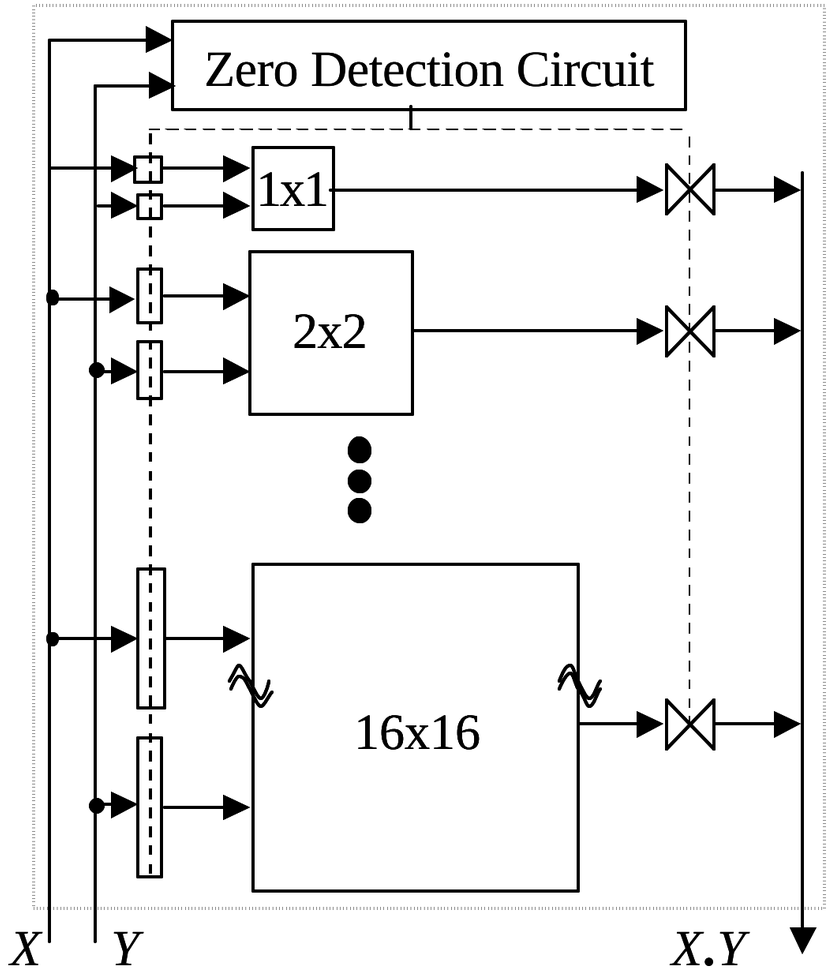

- **What is the optimal ensemble of point systems?**

Formal Statement of the Problem

- **Given:**
  - **Function to be realized (*F*)**
  - **Constraints to be met (*C*)**
  - **A set of point systems (*P*)**
  - **A scenario distribution (*d*)**
- **Form of the solution:**
  - **An ensemble of point systems**
  - **A scenario to point system mapping**
- **Measure of the solution: Power awareness**
- **Problem: Find the solution with the highest measure**
- **Appears to be unsolvable in polynomial time**
- **(Greedy) Heuristics seem to work well**
- **Can be generalized to temporal and spatial-temporal ensembles**

A Near-optimal 4-point Ensemble

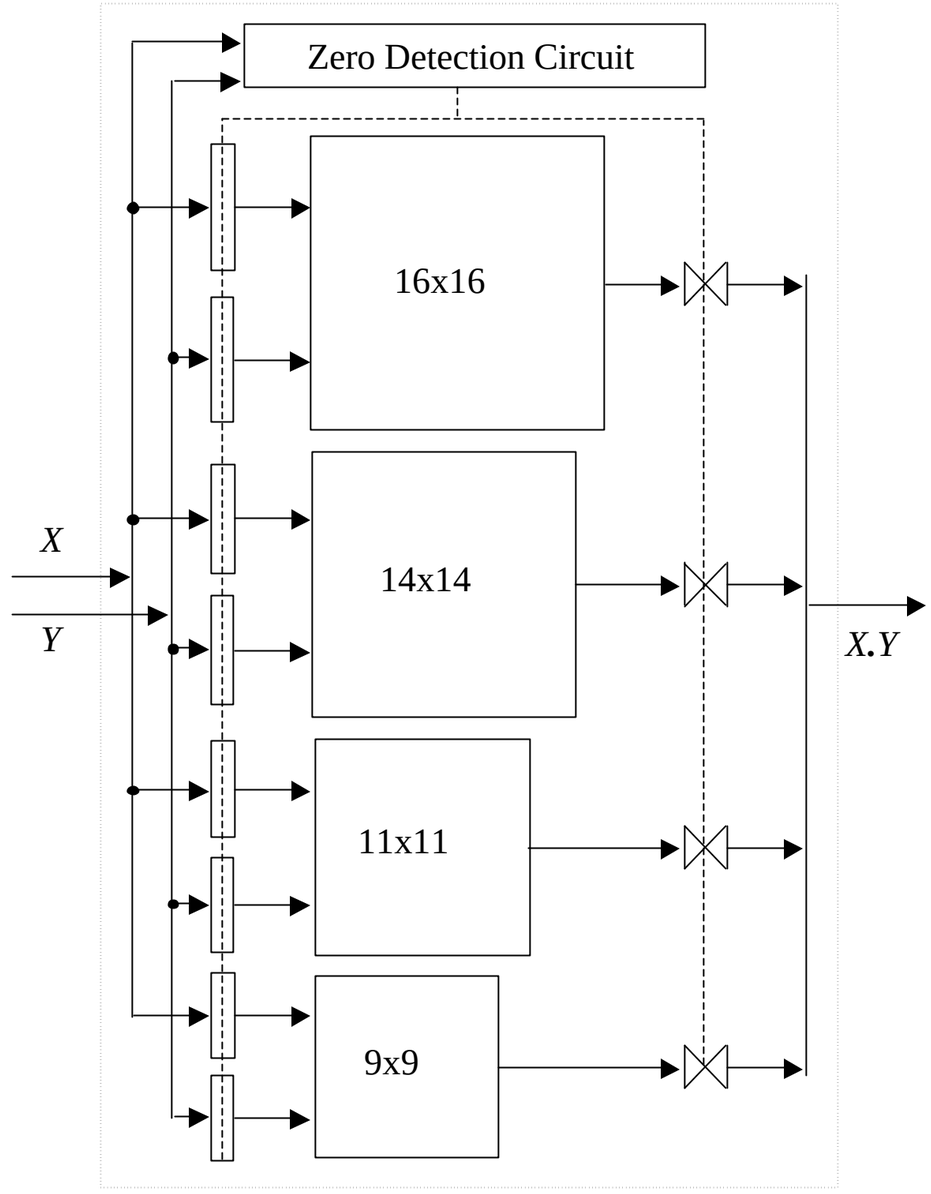

Power-Awareness = 0.92

Power-Aware Register Files

- **Motivation**
  - **Architecture trends point to increasingly energy-hungry files**
  - **Processors typically access only a fraction of registers over typical instruction windows**
  - **Why pay the energy price of full file access?**
- **Objective: Register access energy must scale with the number of registers being accessed over an instruction window**
- **Scenario: Number of distinct registers accessed in an instruction window of specified length**
- **Available point systems: 1, 2, 4, 8 … word register files**

Scenario Distributions

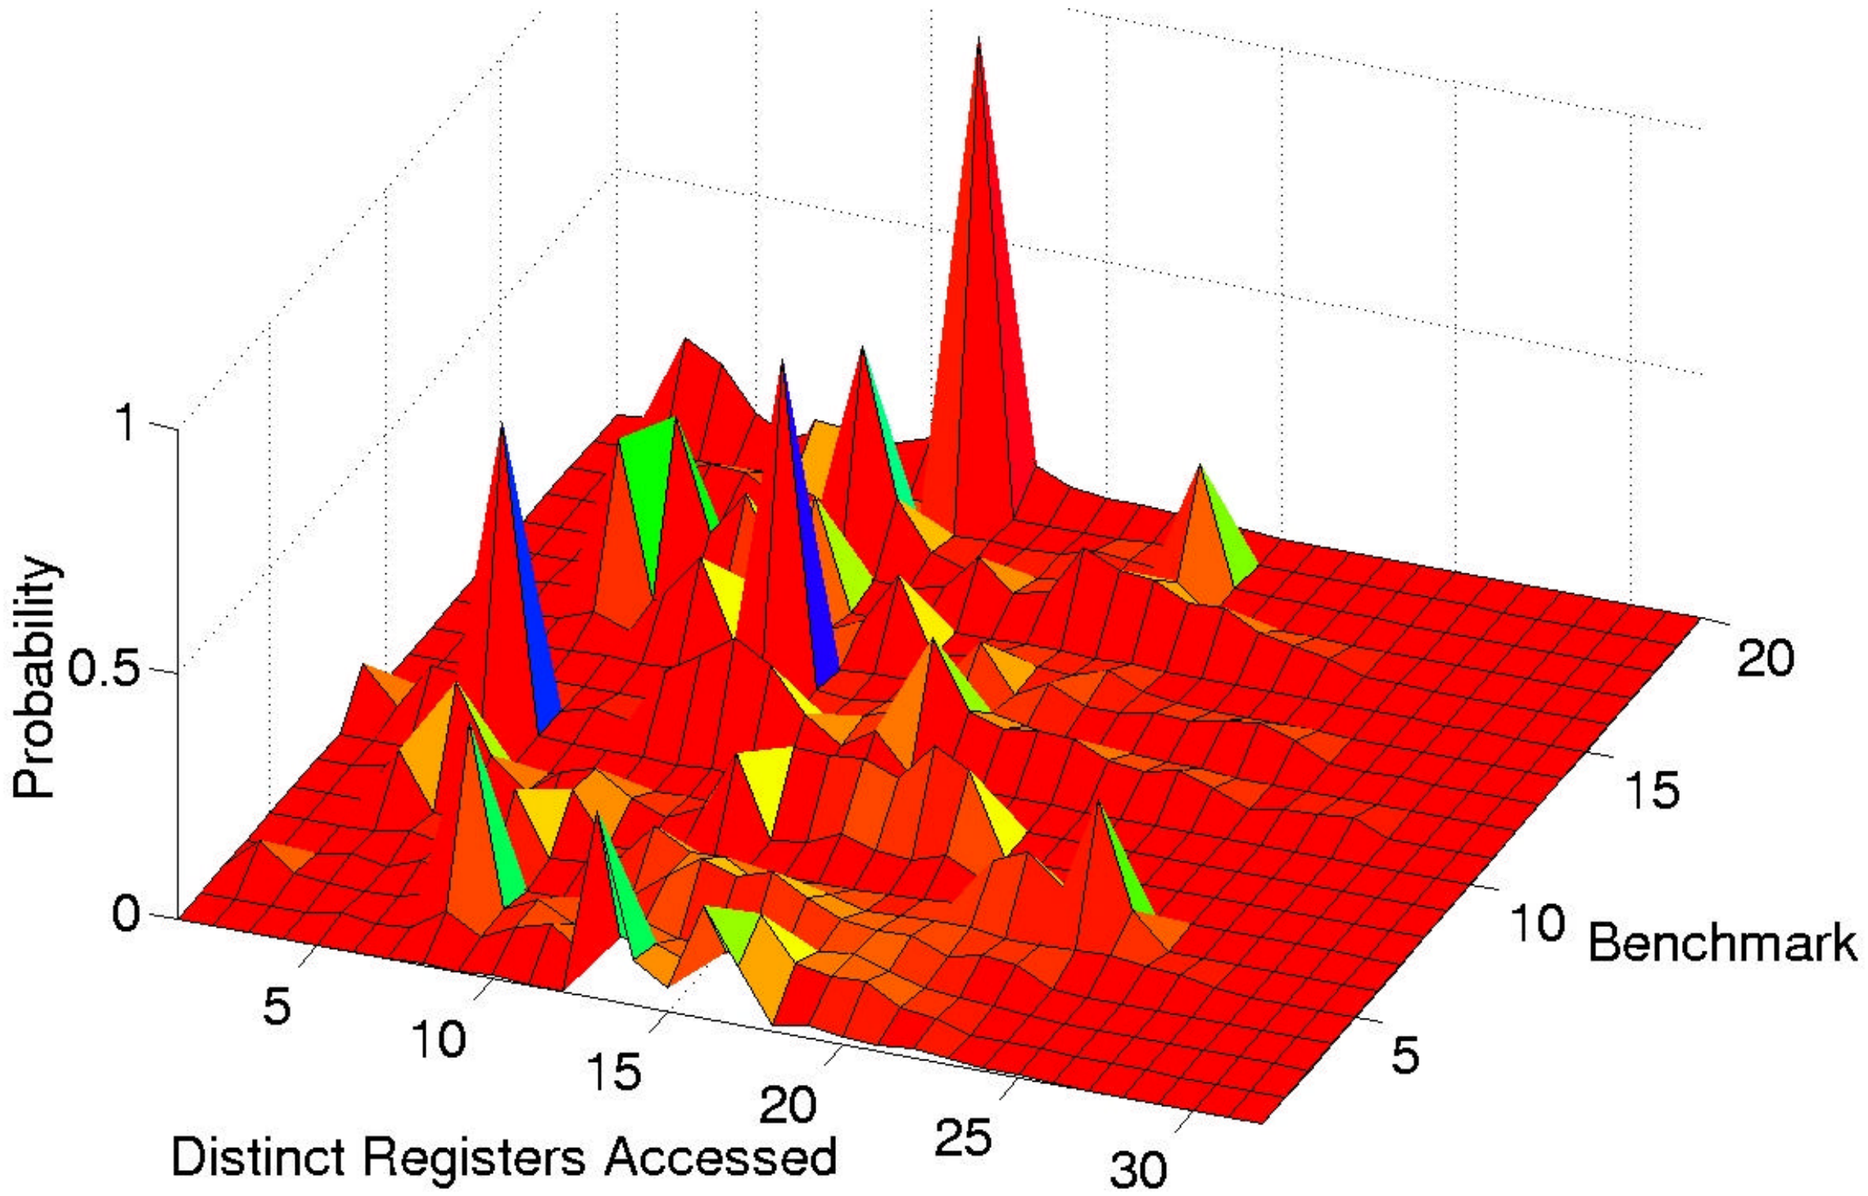

>70% of the time, <16 registers accessed in a 60 instruction window

Window Locality

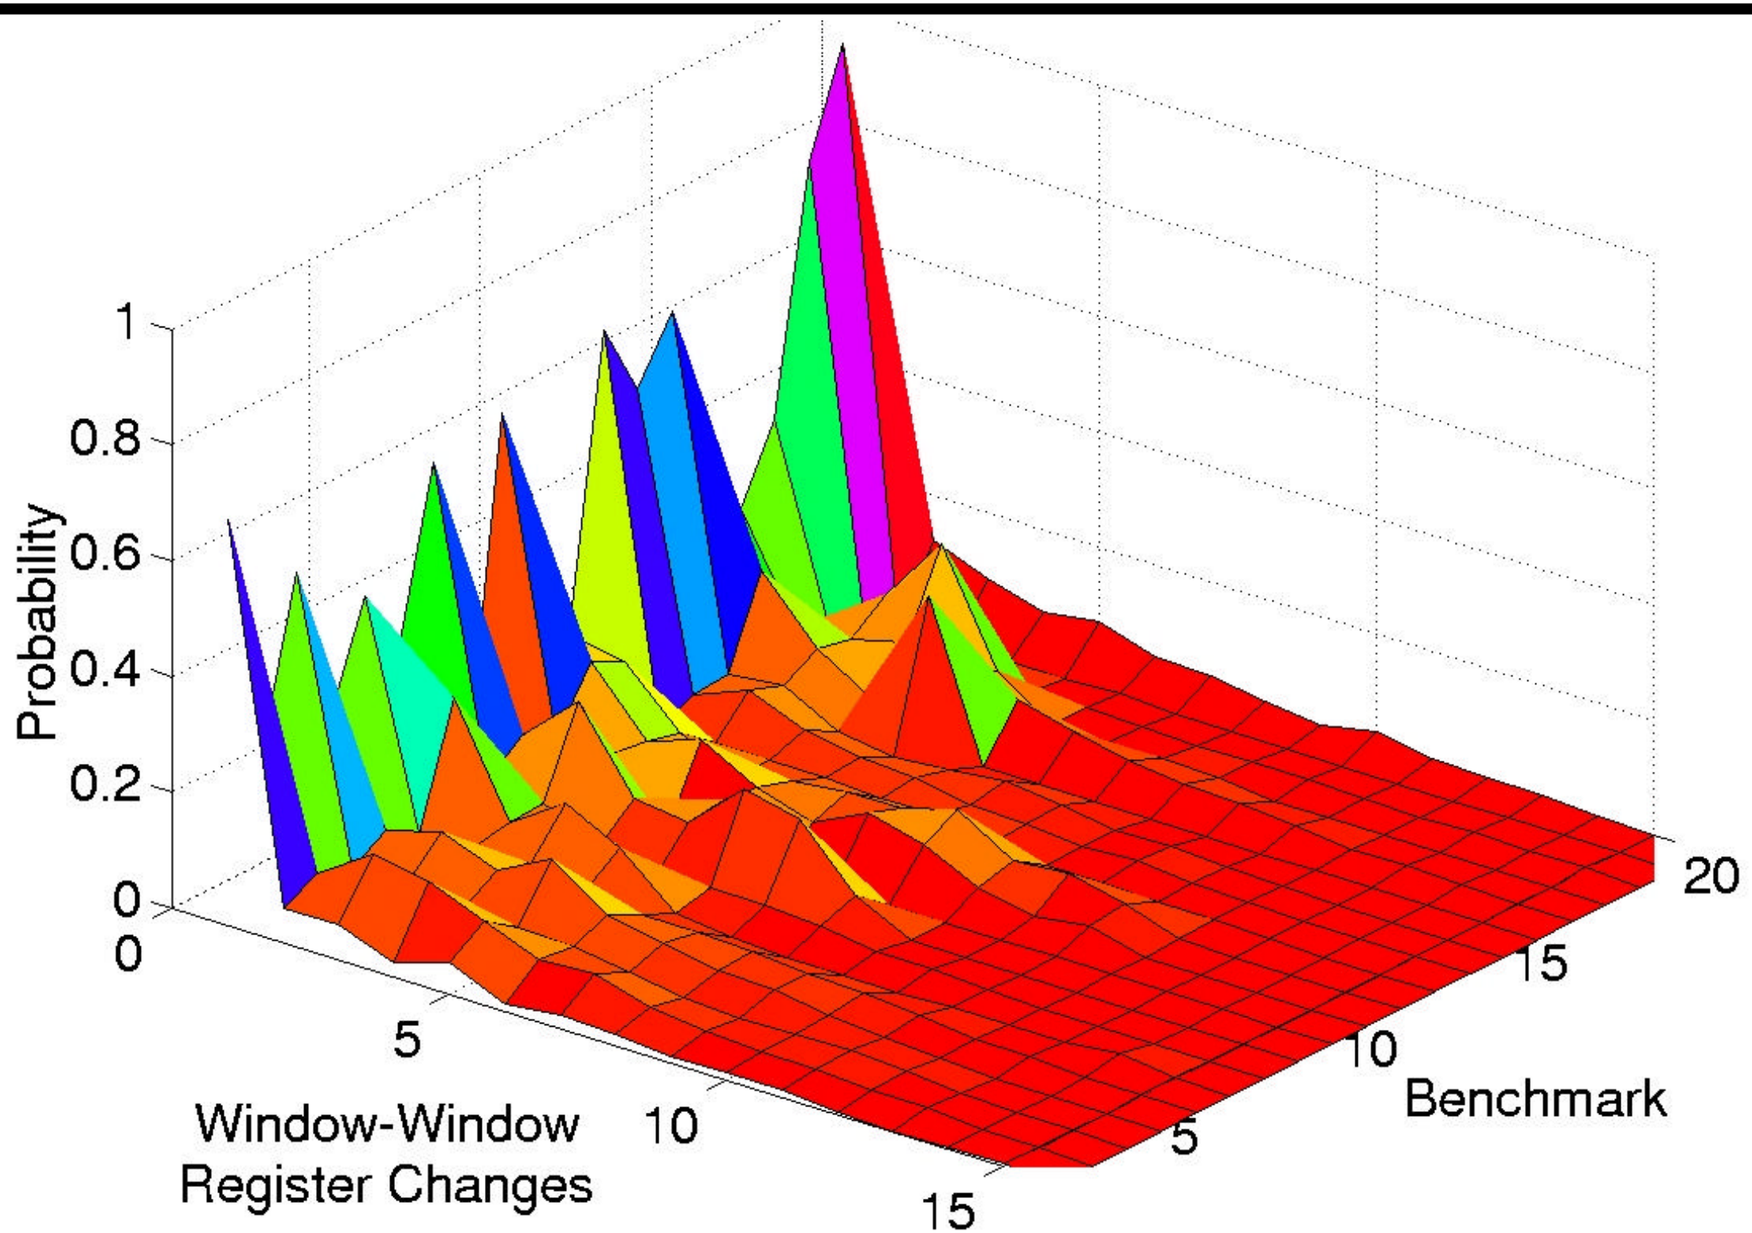

>85% of the time, <5 registers change from window to window

Candidates

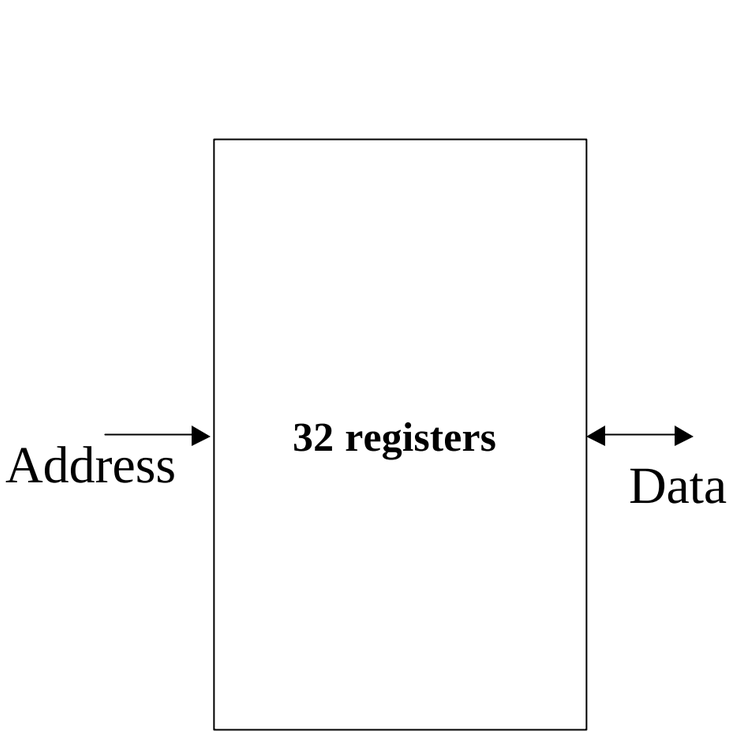

Monolithic File

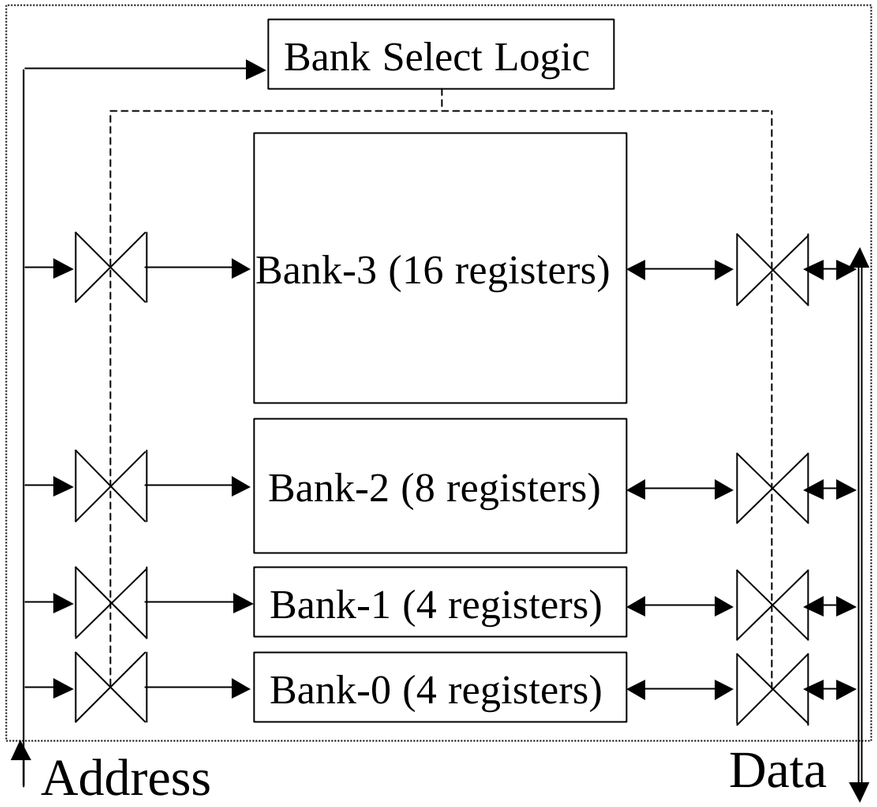

Segmented File

Power-Awareness Comparisons

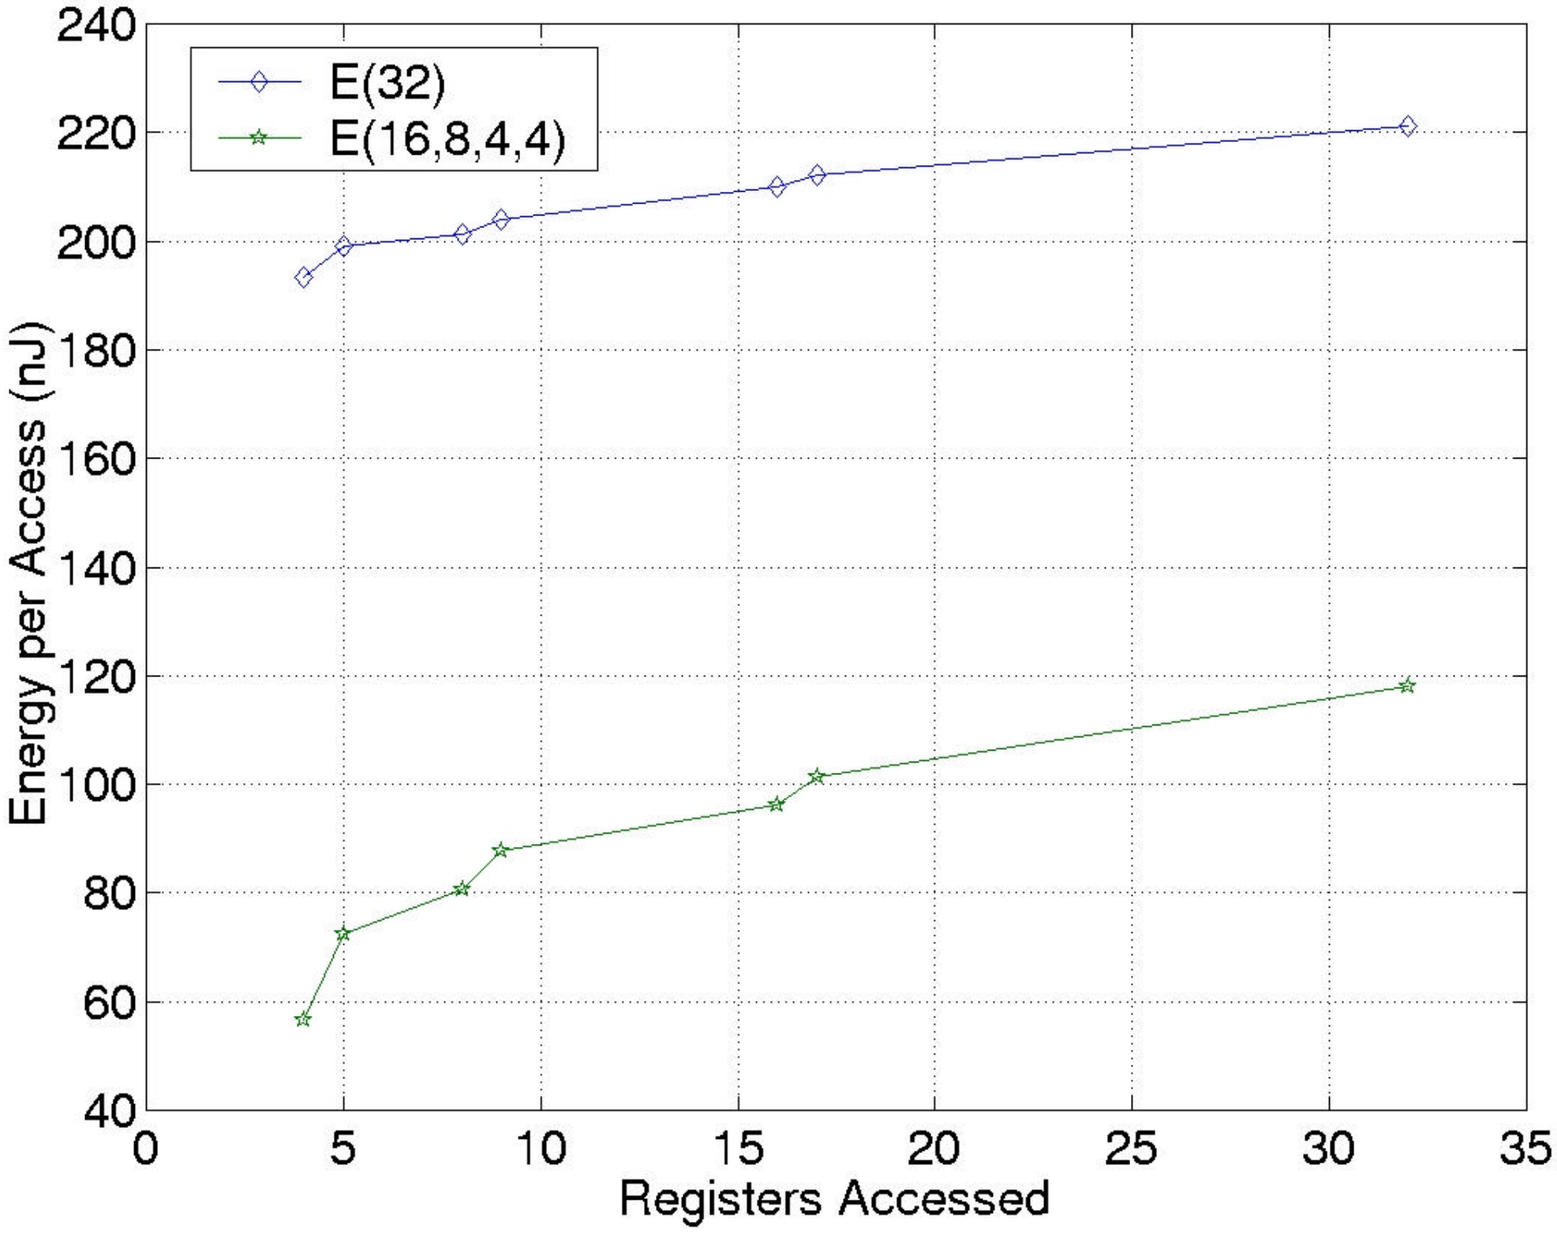

Power-Awareness Increases by 2-3x

Power-Aware Digital Filters

- **Motivation:**
  - Adaptive filters used in communications applications dissipate significant energy
  - Filtering requirements change with desired quality and channel conditions
  - Why run the filter at maximum precision and taps?
- **Objective: Energy consumed by a filter must scale with the word-length precision and taps**
  - **Scenarios: <Desired Taps, Desired Precision>**
  - **Point systems: All <*m* taps, *n* bits> filters**

Scenario Distribution

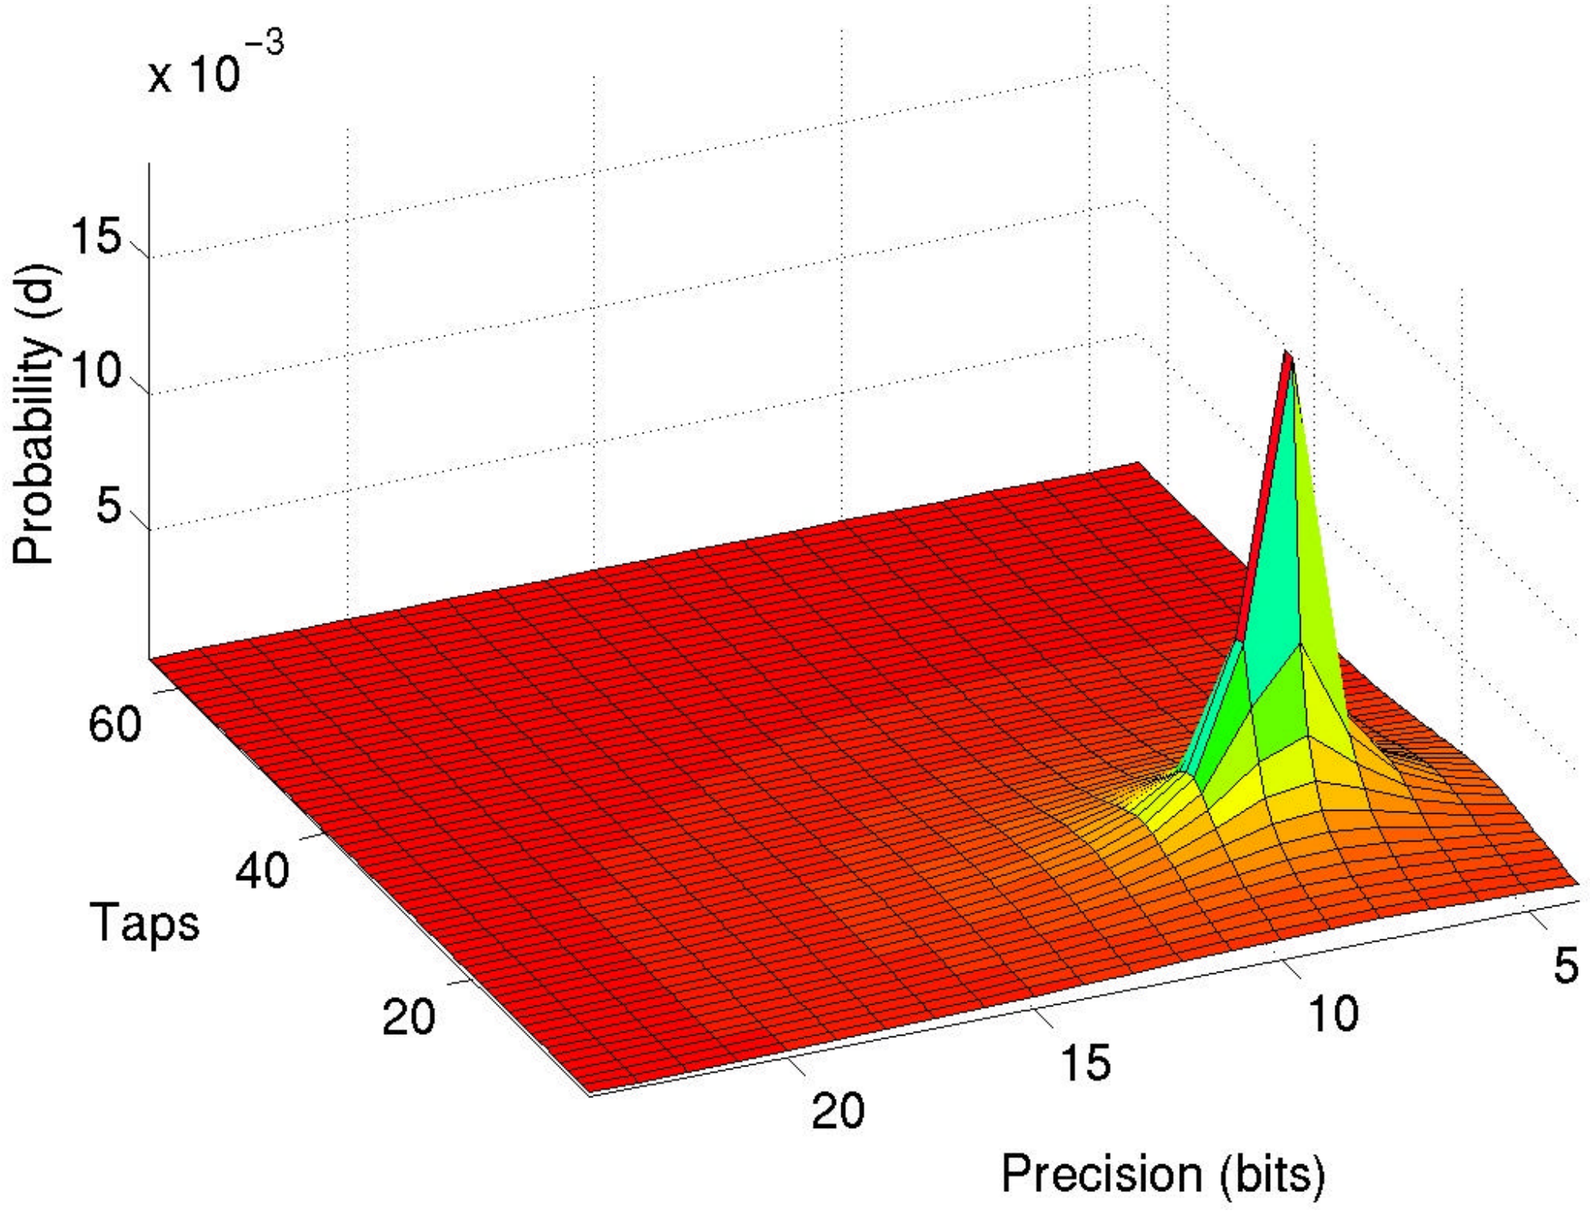

Candidates

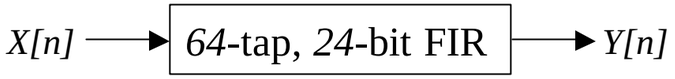

Monolithic Filter

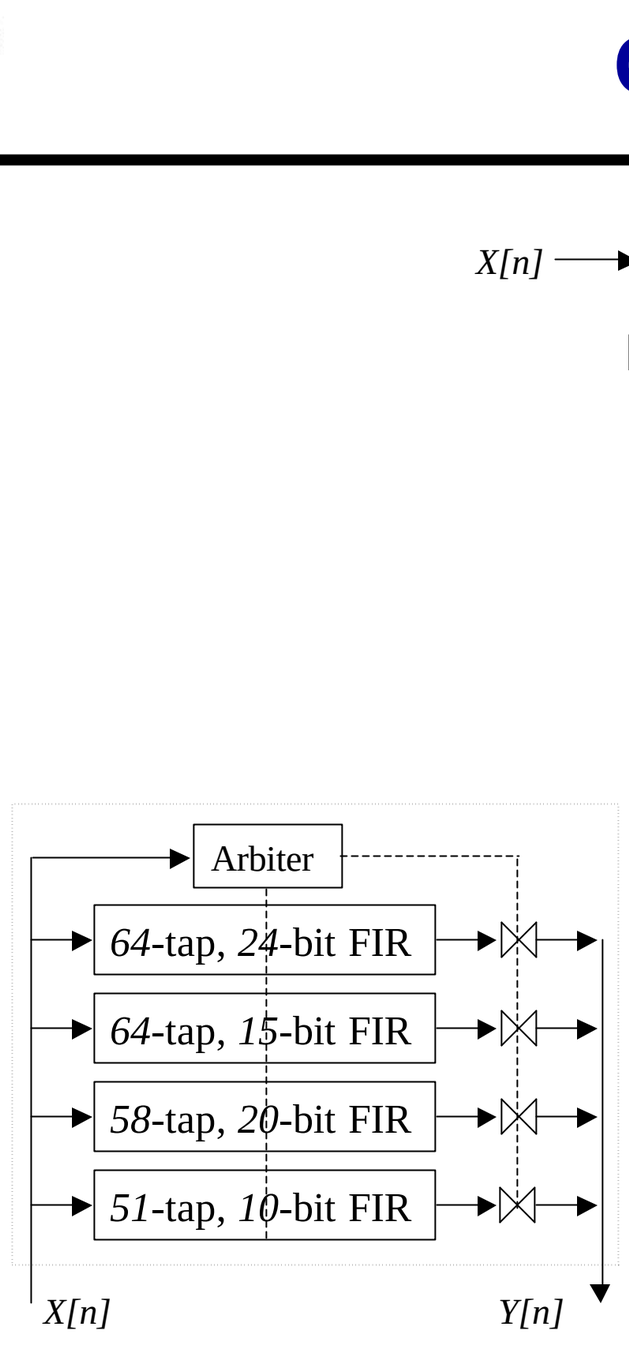

Optimal 4-point Ensemble

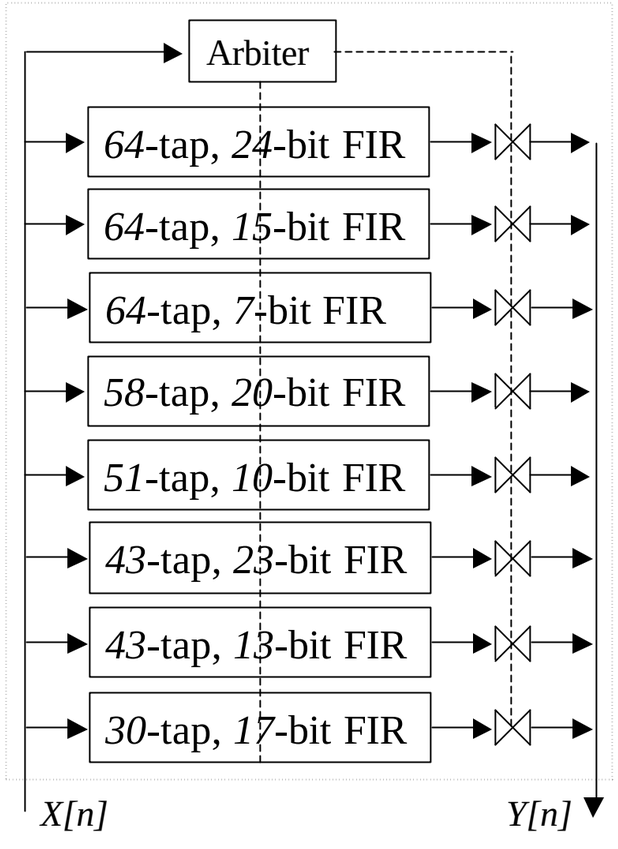

Optimal 8-point Ensemble

Monolithic Filter

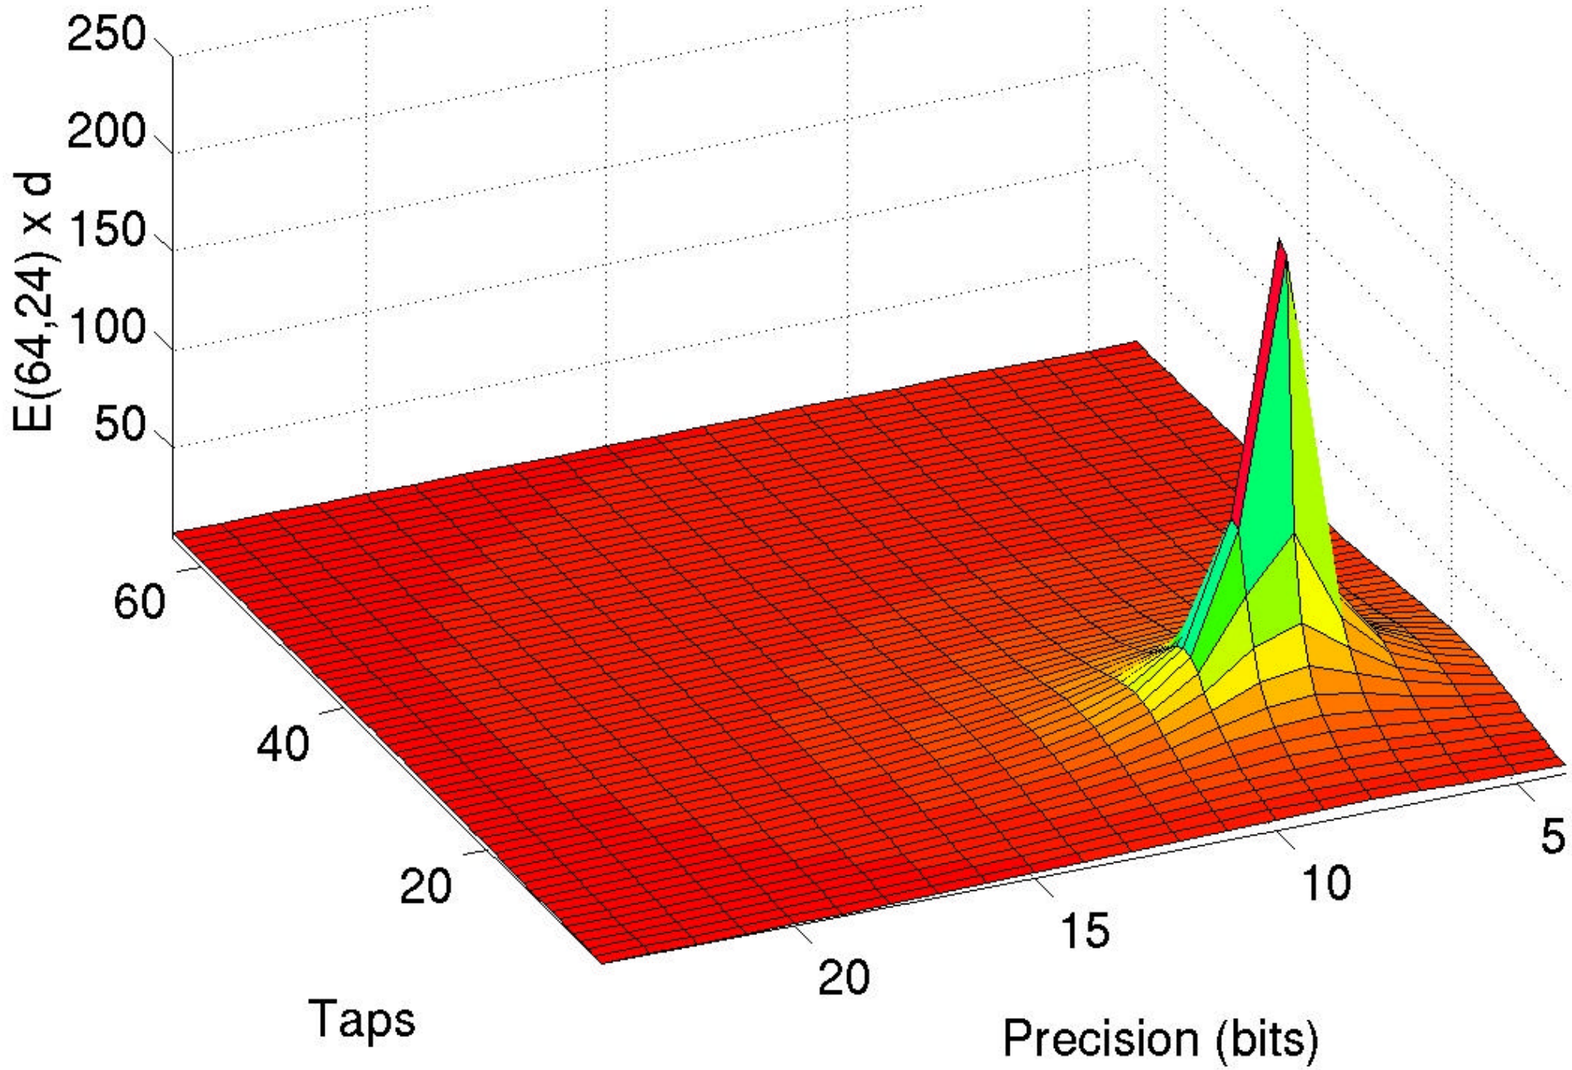

Power-Awareness = 0.51

4-point Ensemble

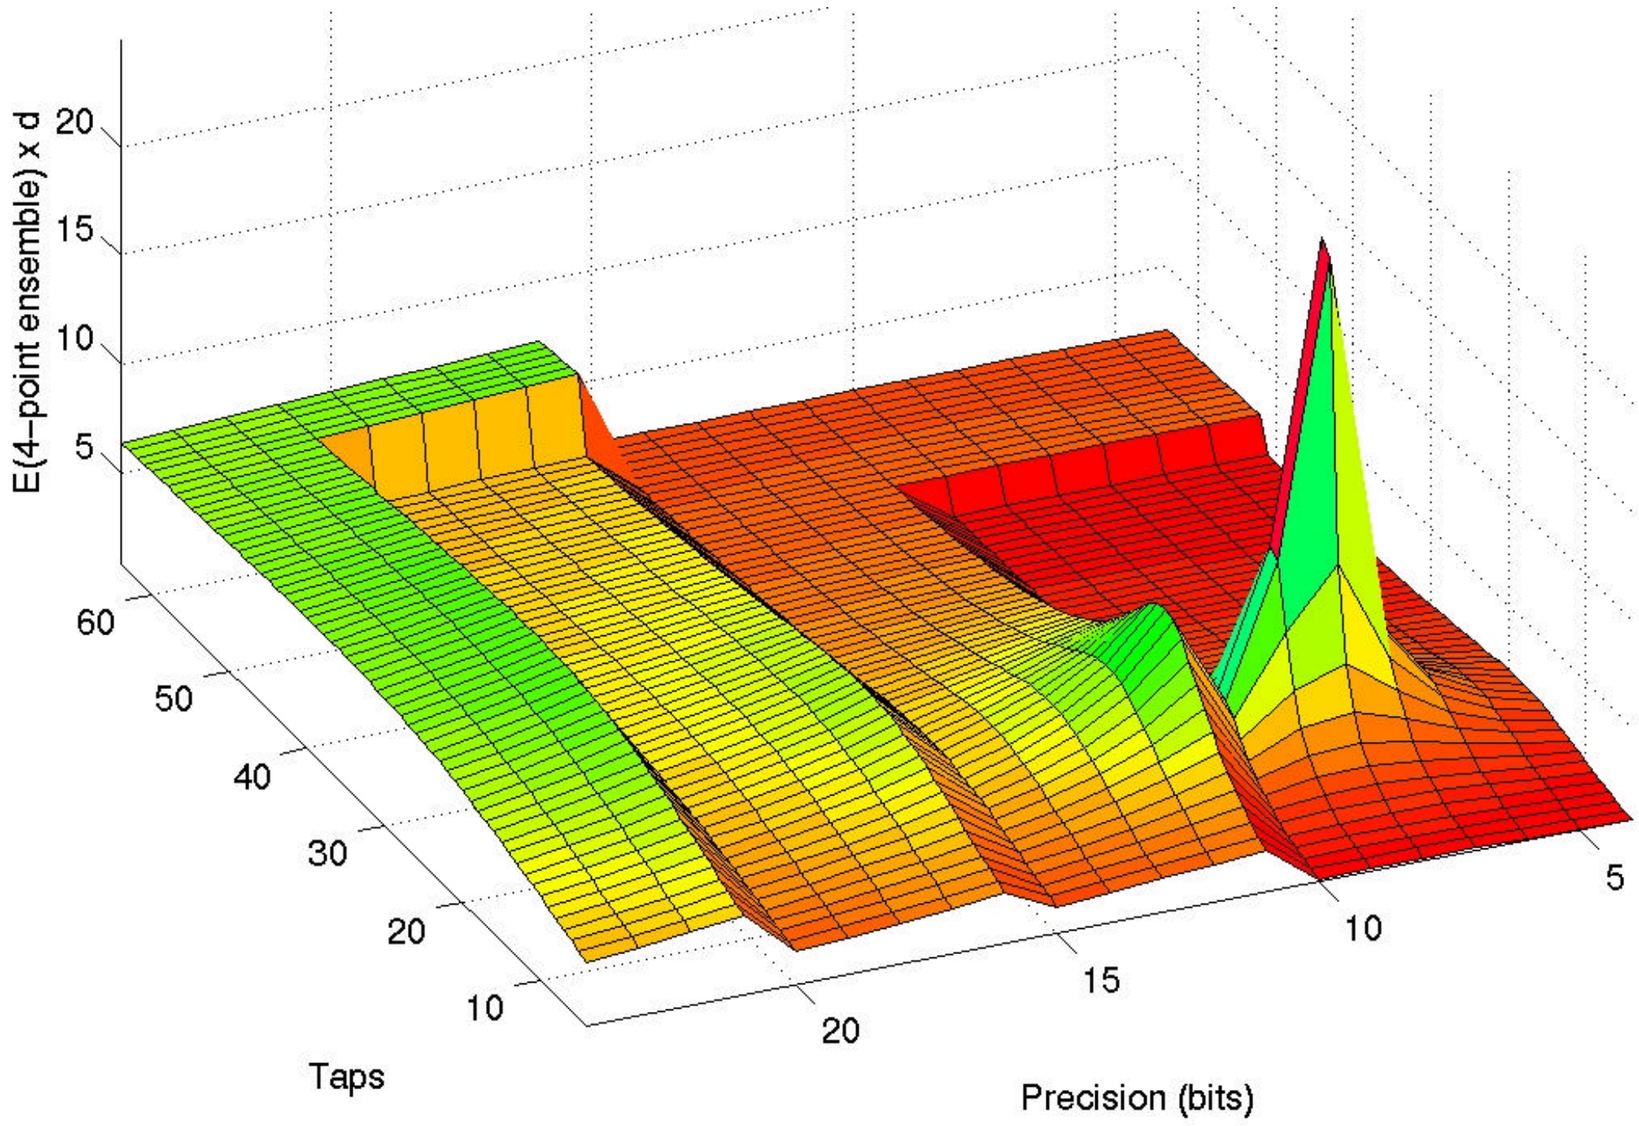

Power-Awareness = 0.82

8-point Ensemble

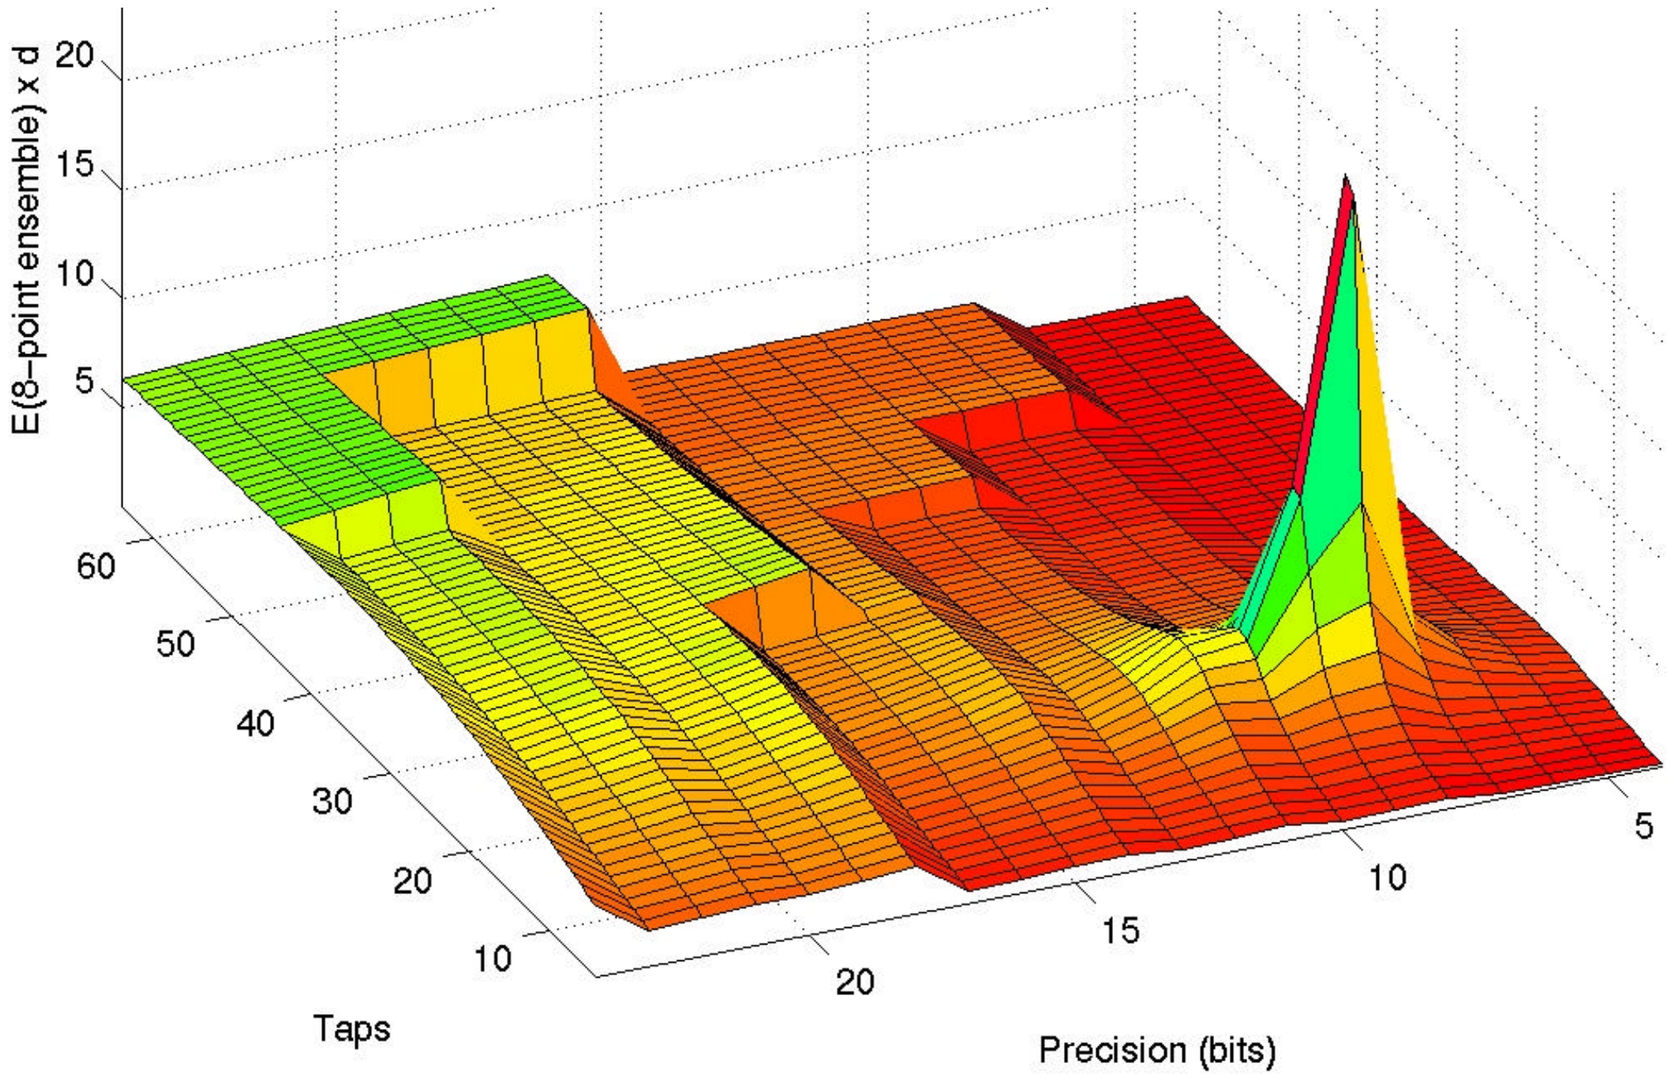

Power-Awareness = 0.90

Perfect System

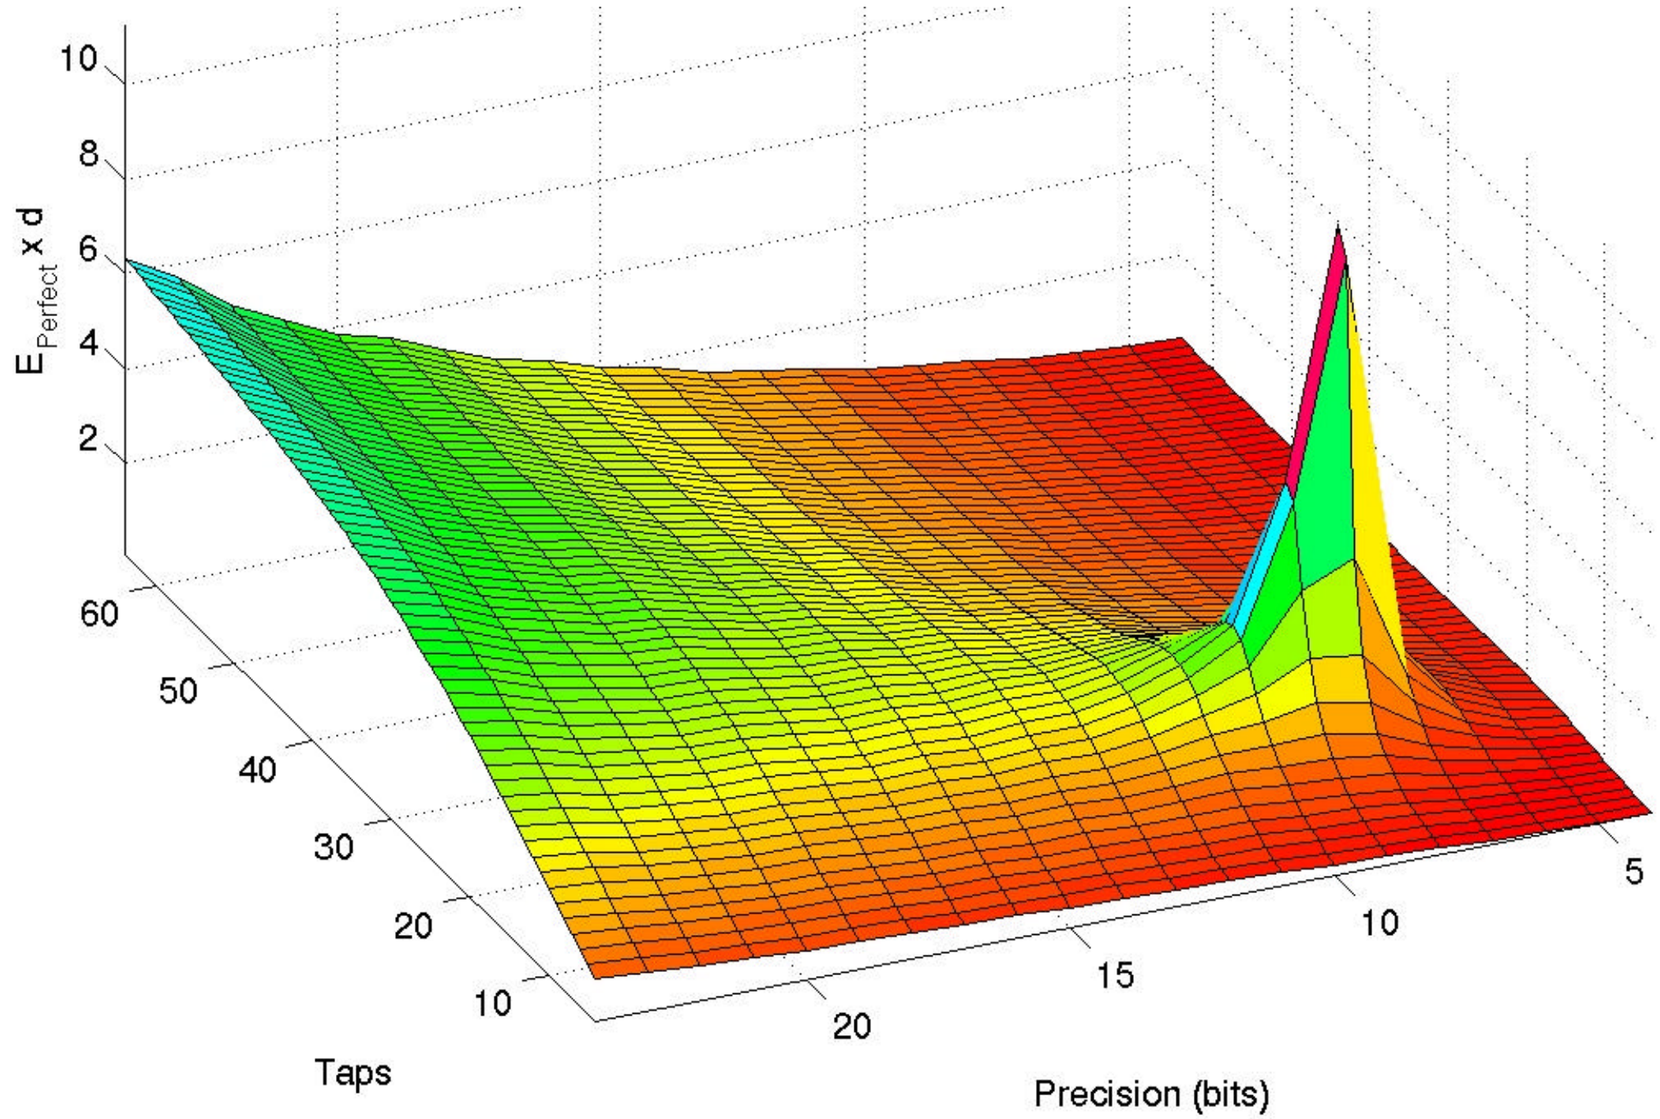

Power-Awareness = 1.0

Power-Aware Processors

- **Motivation:**
  - **Processor workloads vary significantly**
  - **Tremendous energy savings by spreading workload to occupy all available time (by lowering V_{dd} and operating frequency)**
  - **Why pay the energy price of a full workload?**
- **Objective: Energy consumed by a processor should scale with its workload requirement**
- **Scenarios: Workload ($\hat{I}$ [0,1])**
- **Point systems: Processors with V_{dd}, frequency customized for a workload**

Candidates

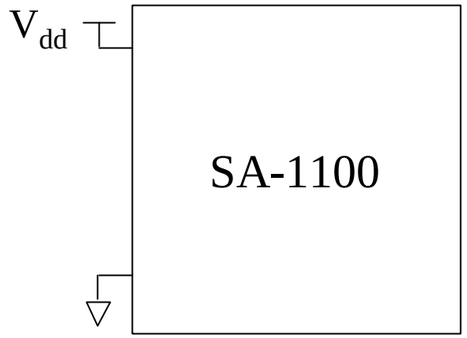

Fixed Voltage Processor

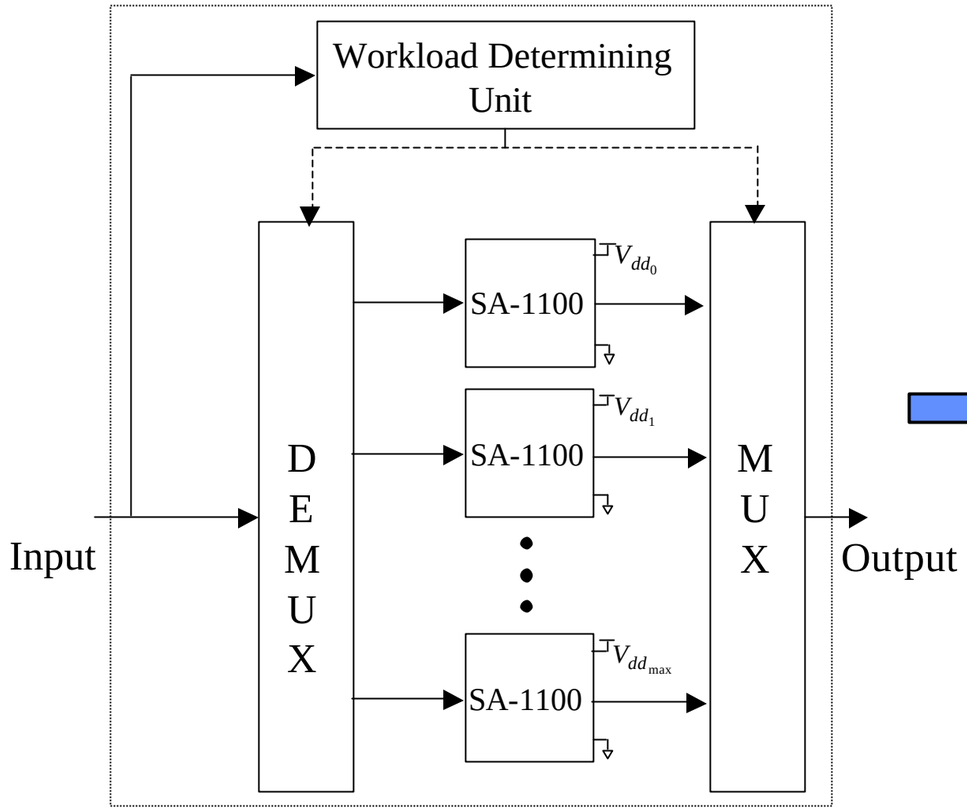

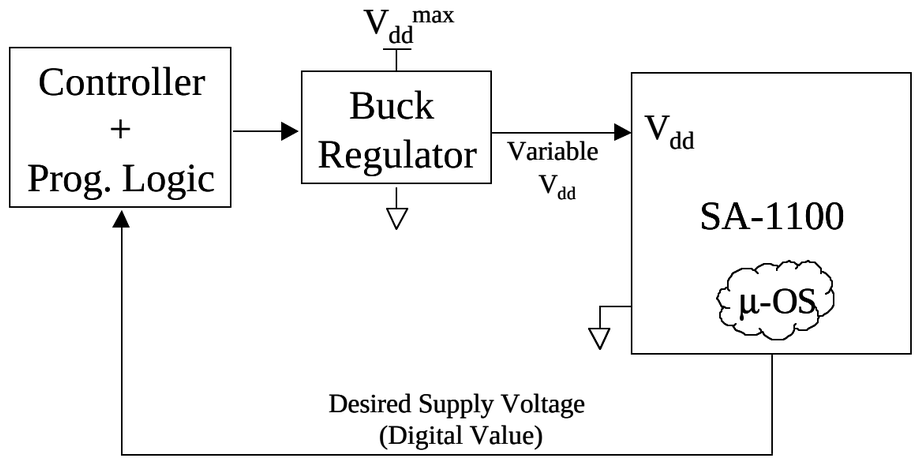

Dynamic Voltage Processor

Power-Awareness Comparisons

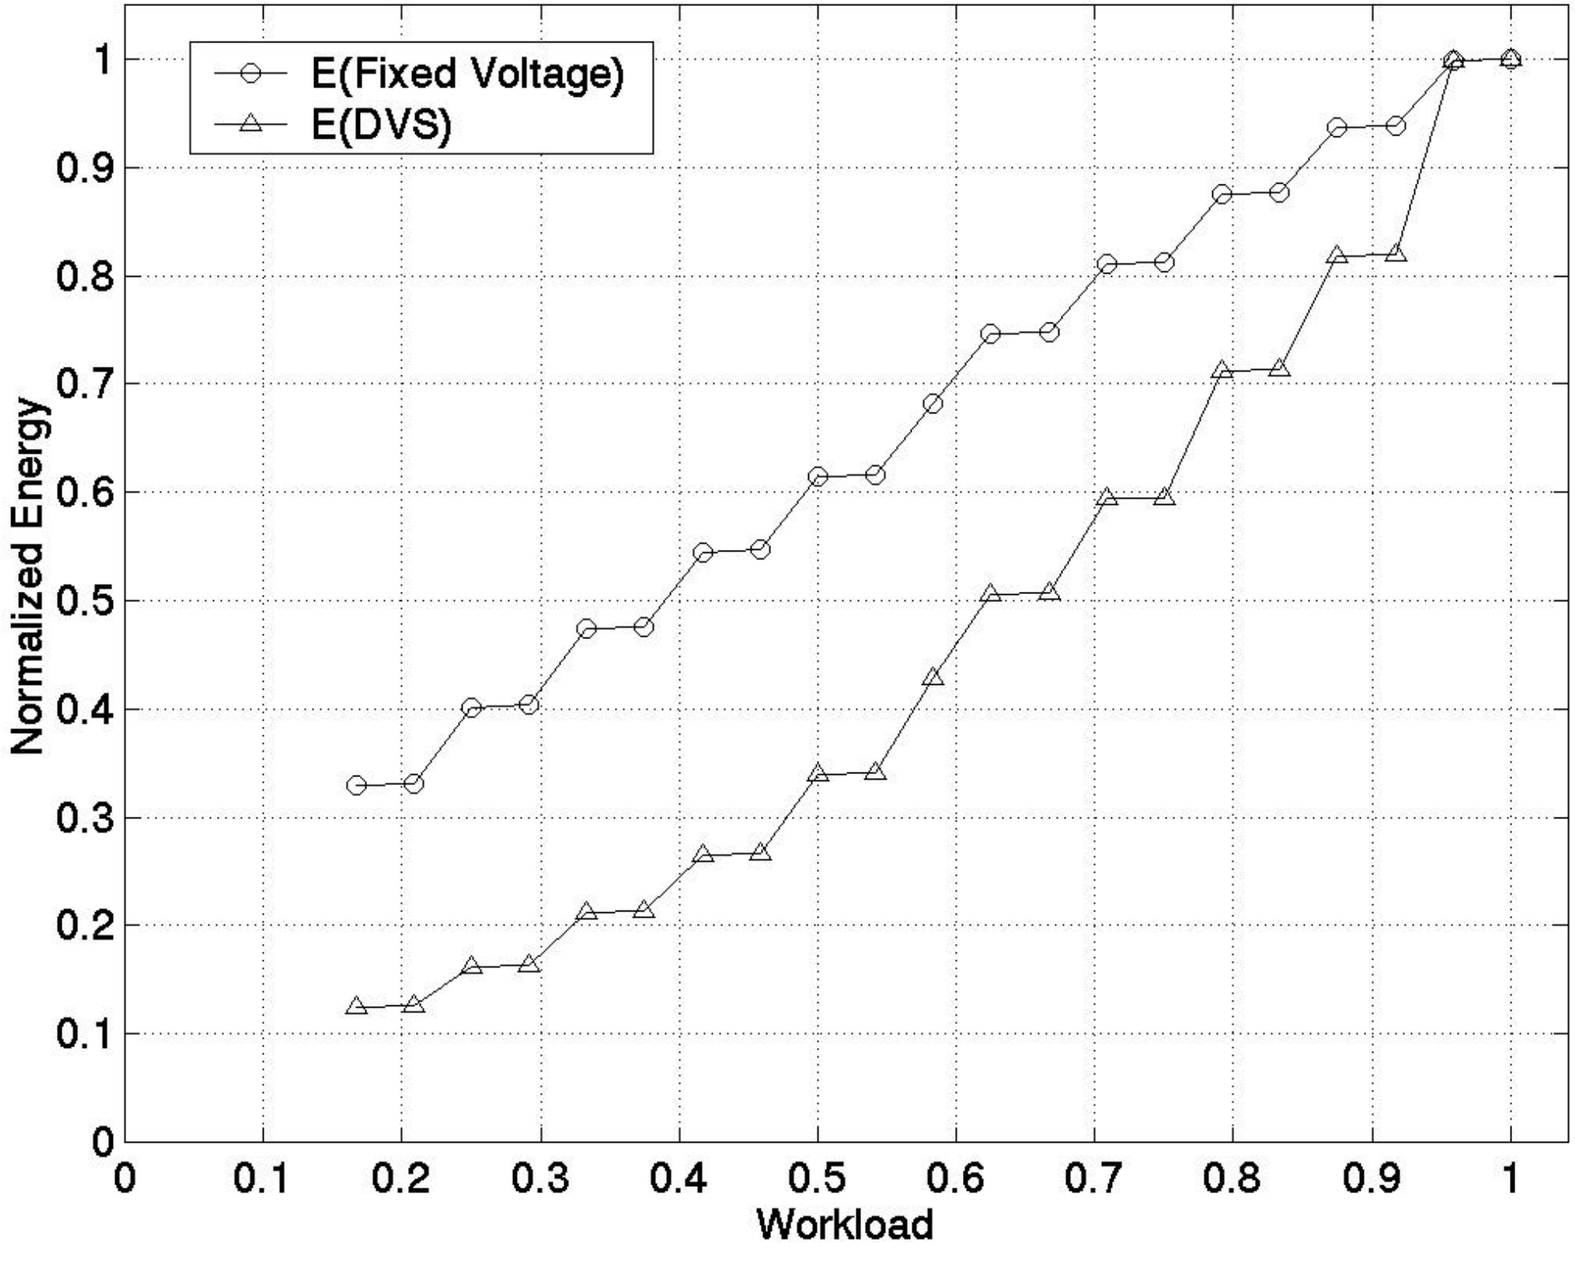

DVS 1.6x more power-aware than fixed-voltage system

Analog-Digital Converters

Contributed by Kush Gulati, MIT [ISSCC'01]

- **Motivation:**
  - A/Ds have non-trivial system-level power-budgets
  - User/algorithms might be able to tolerate low quality (resolution)
  - Signal statistics might allow variable sampling rates
- **Objective: Conversion energy must scale with the desired sampling rate and resolution**
- **Scenarios: <Rate, Resolution>**
- **Point systems: All <Rate, Resolution> converters**

Candidates

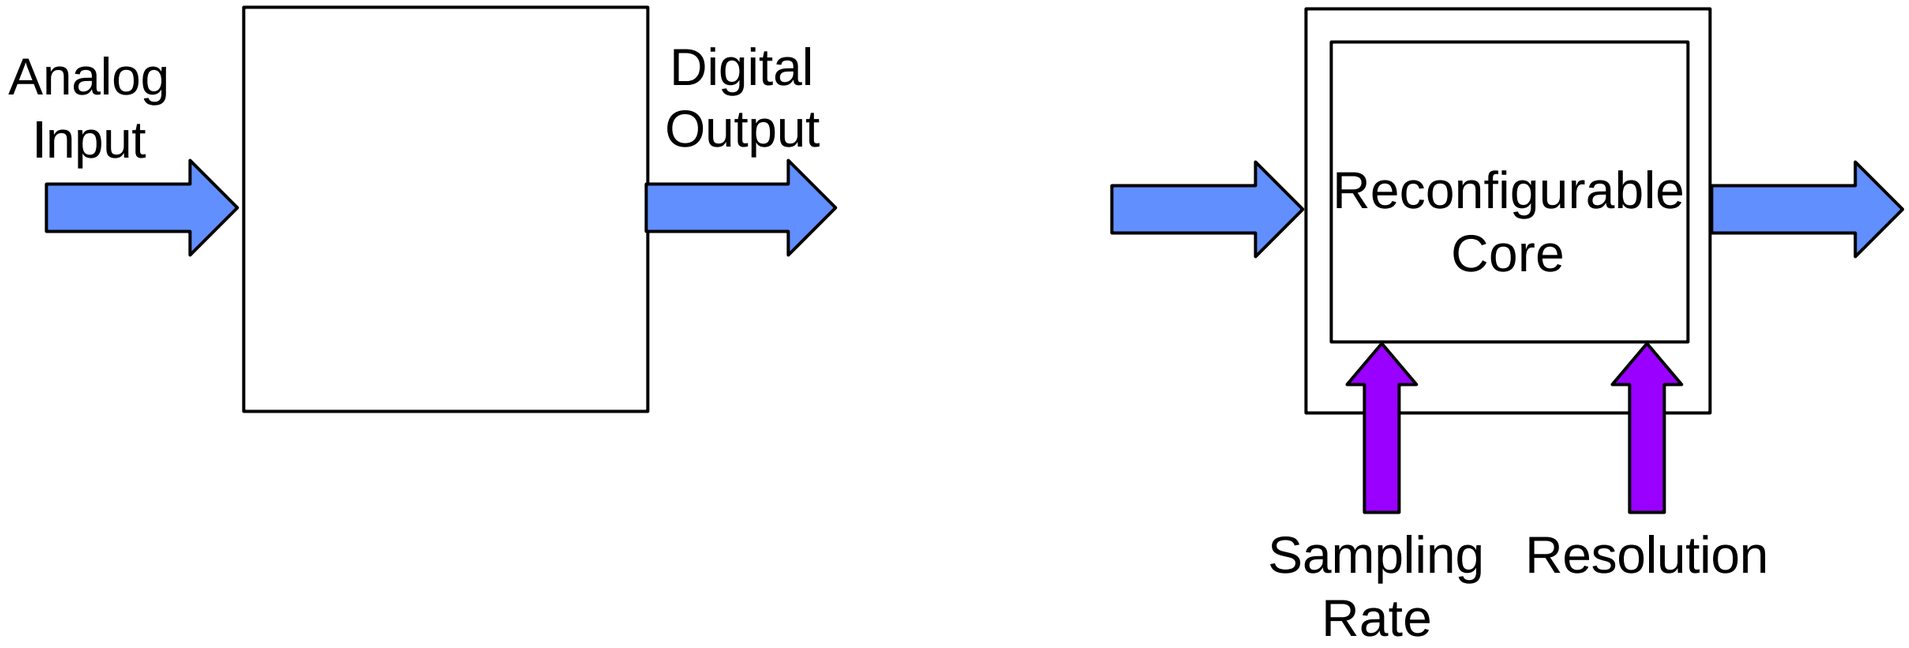

Conventional A/D

Power-aware A/D

Scenario Diversity in A/Ds

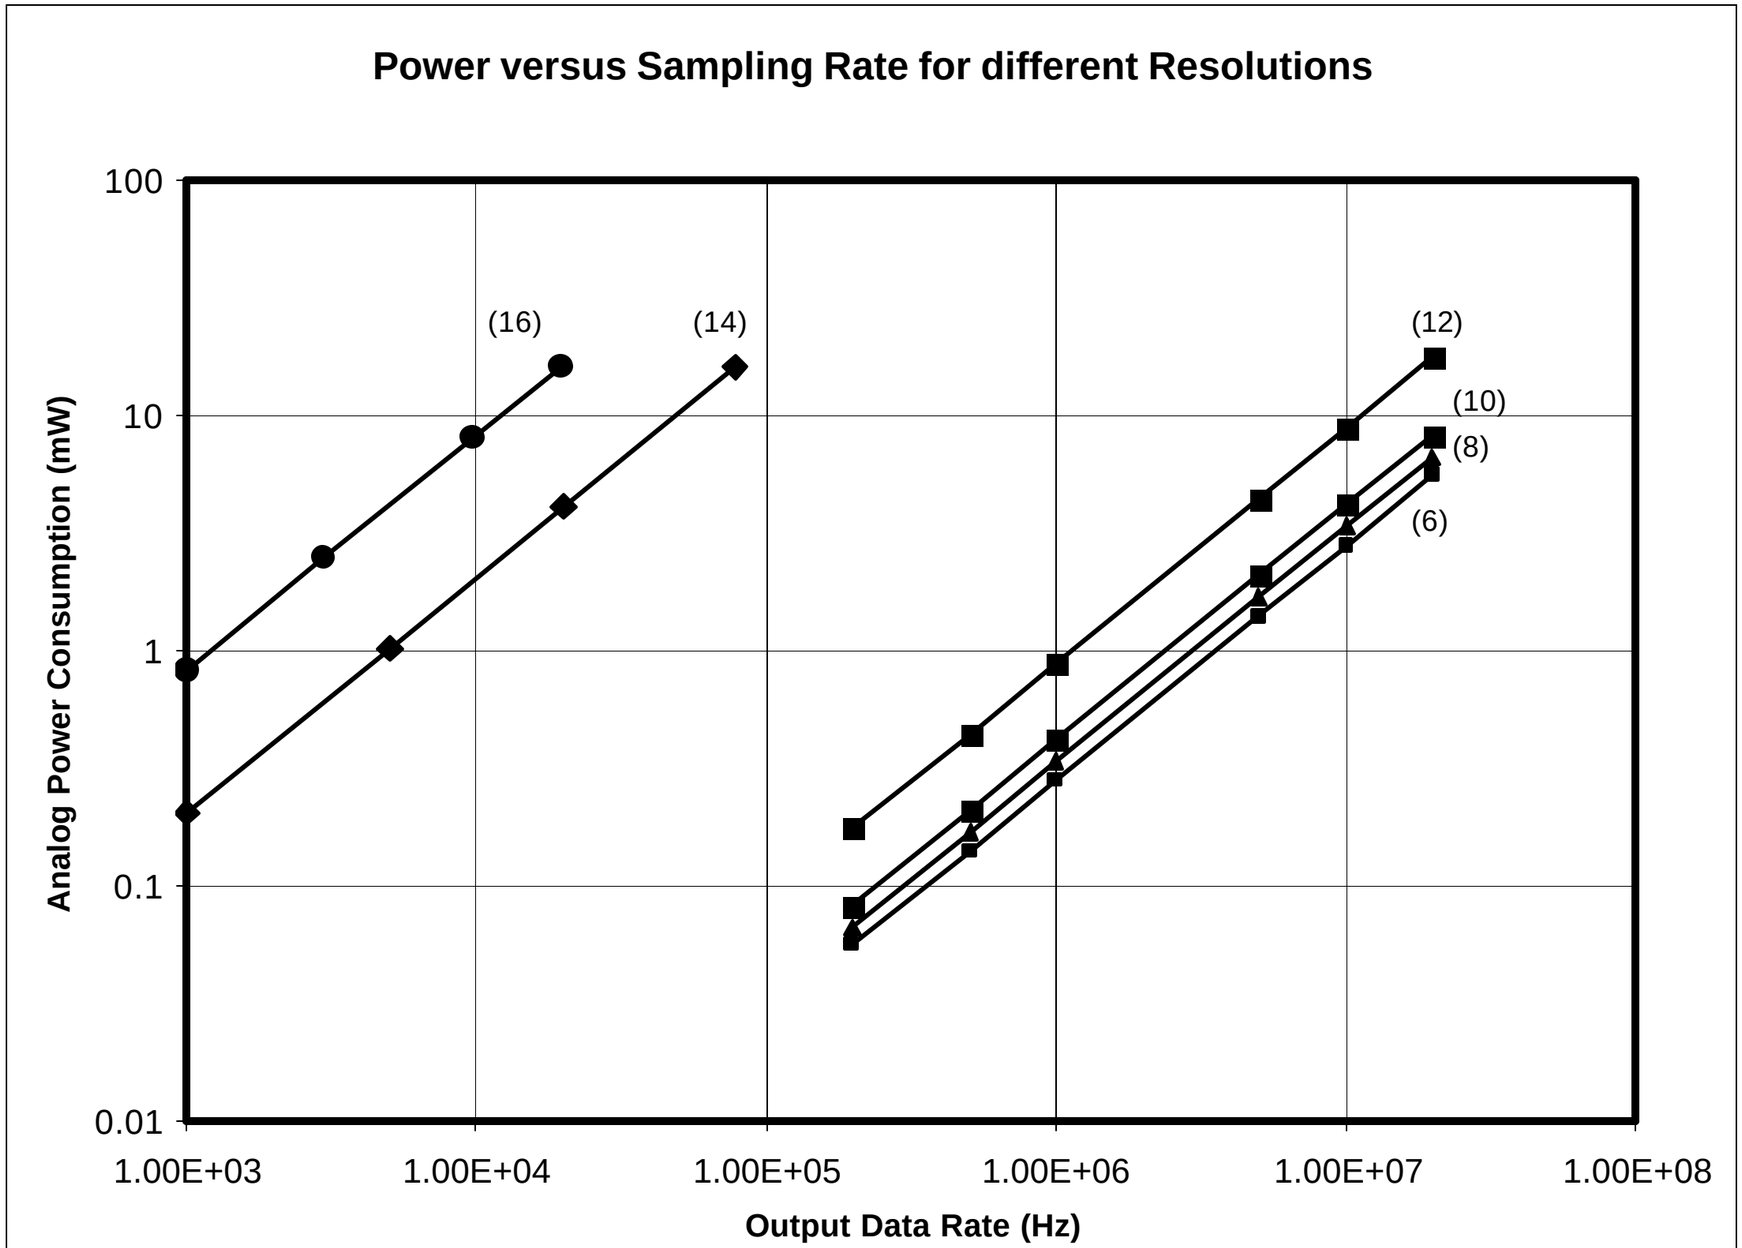

Power-Awareness Comparison

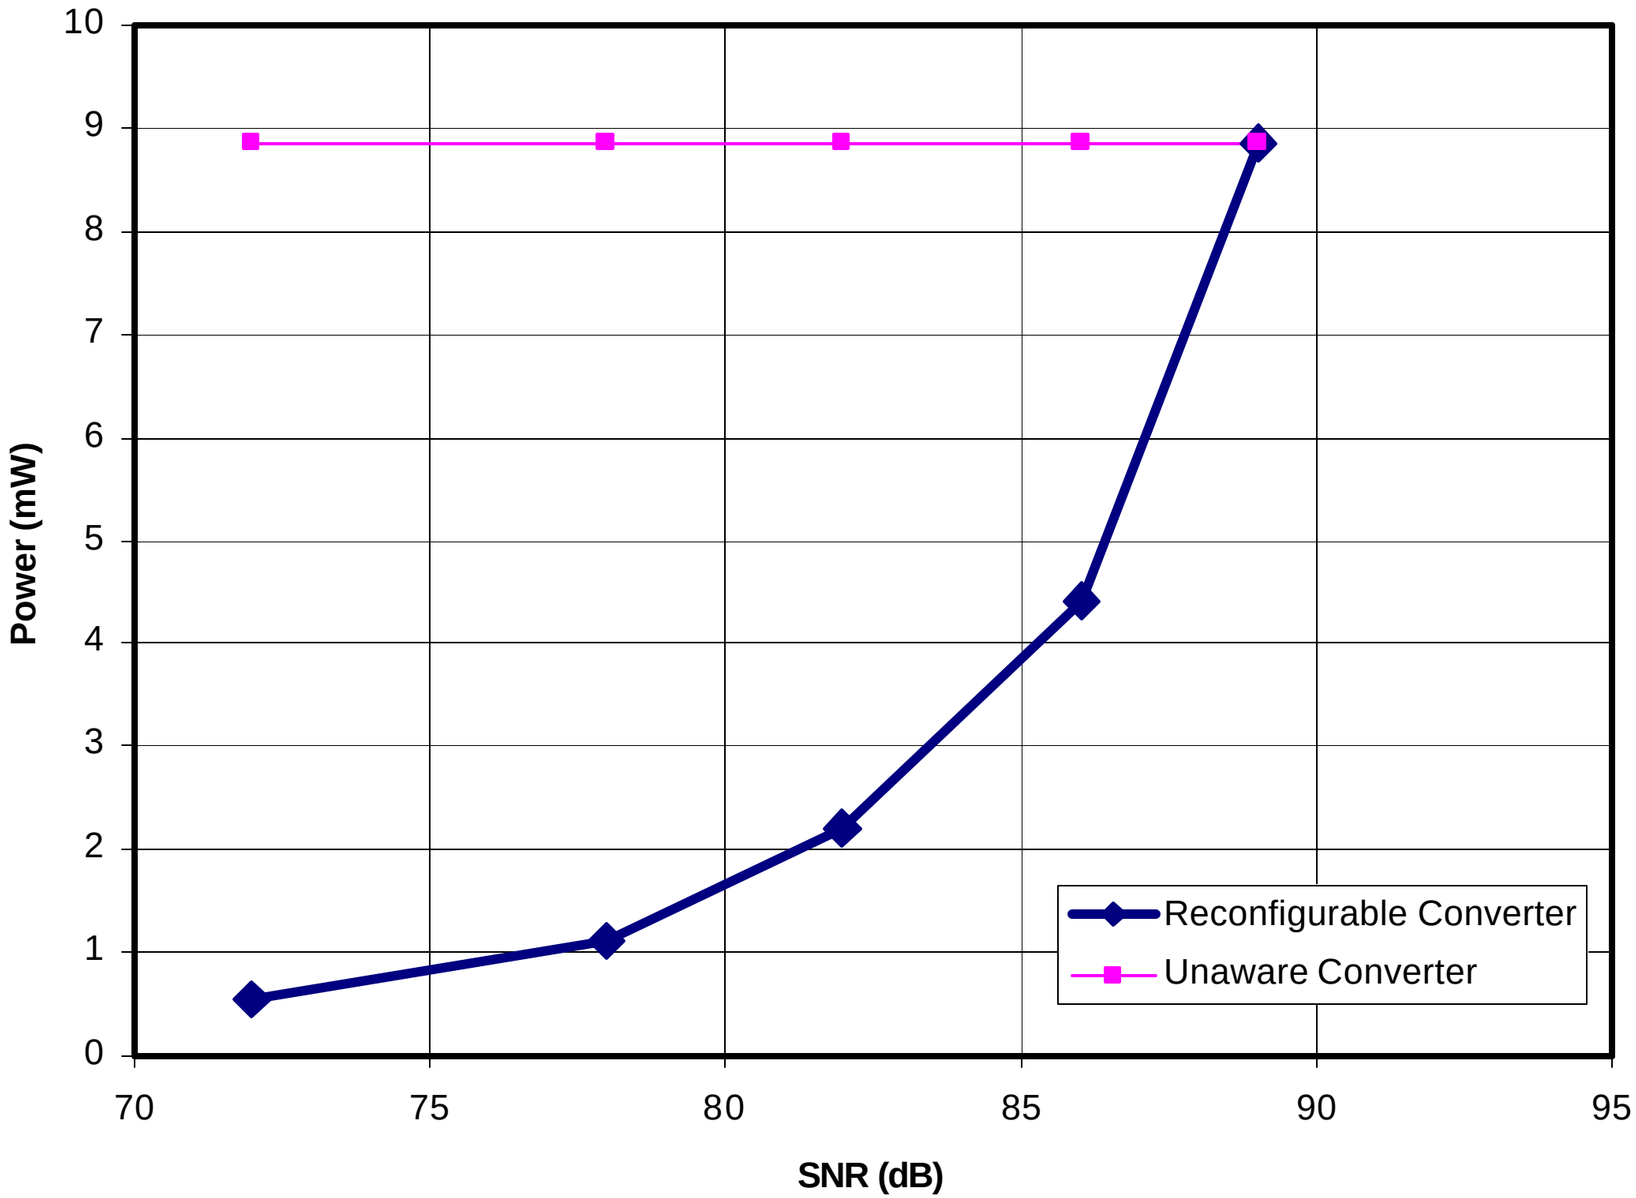

Power-Awareness increases from 0.31 to 0.81

Wireless Data-Gathering Networks

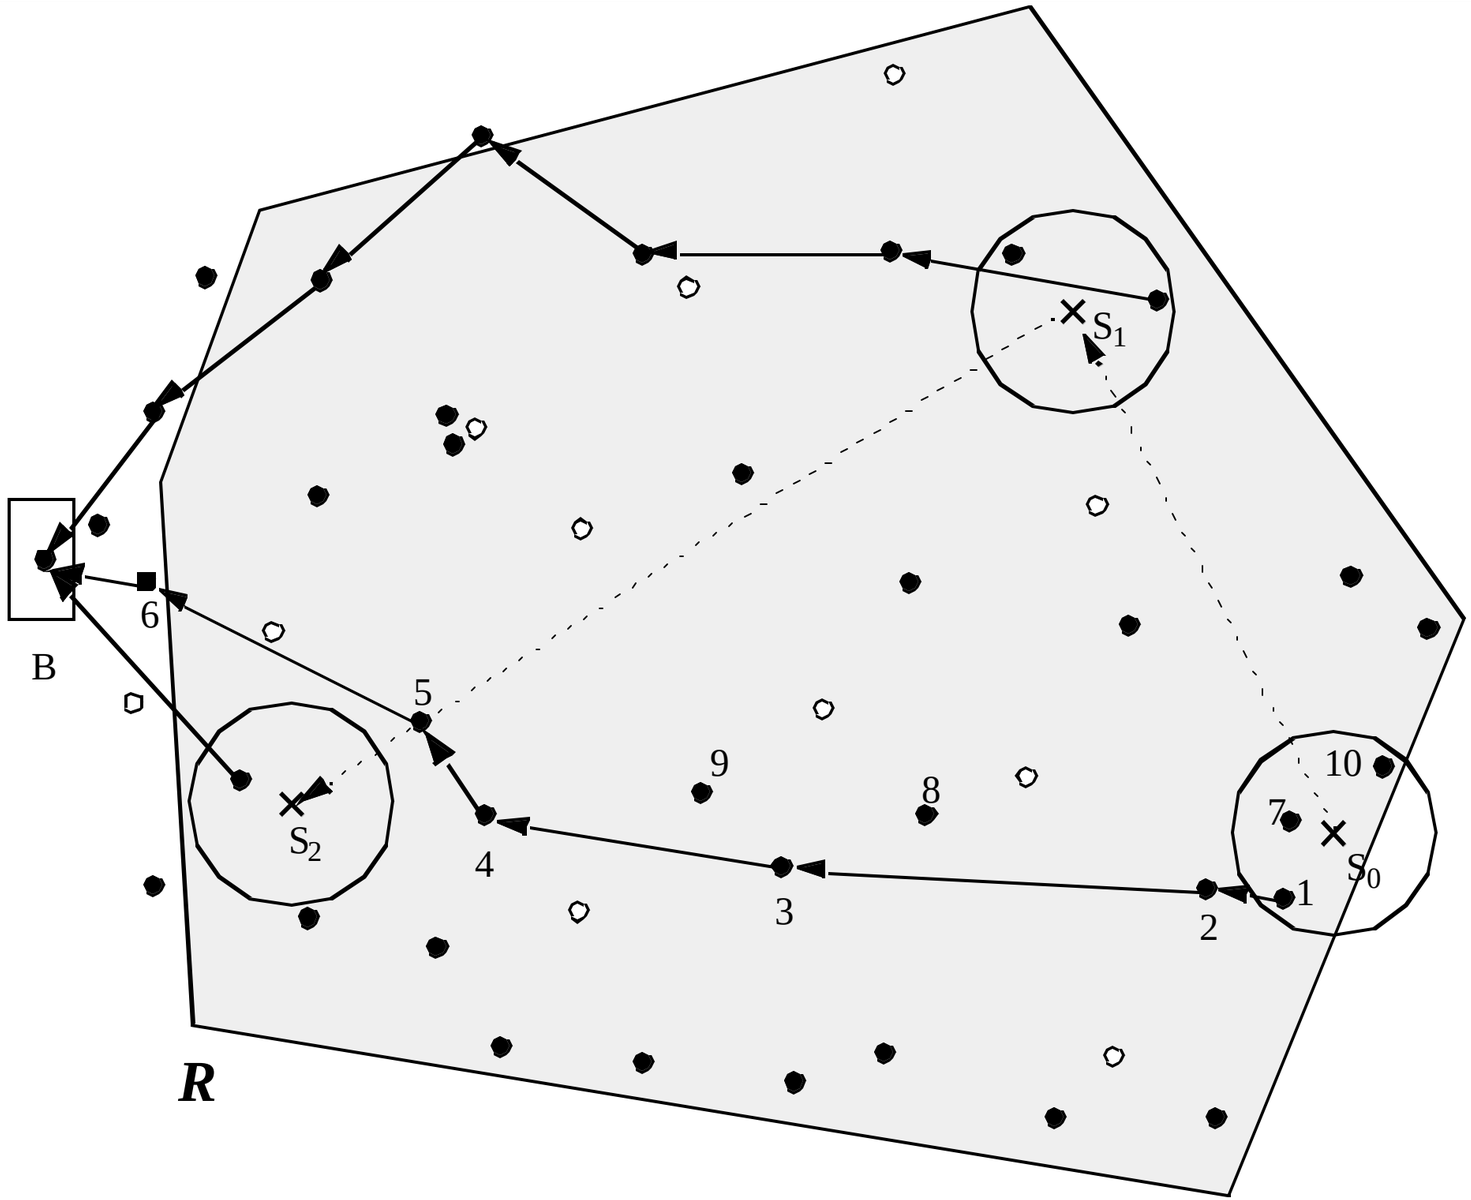

- **Energy constrained nodes deployed to observe a source in a specified region**

Power-Aware Wireless Networks

- **Motivation:**
  - **Key challenge in data-gathering networks is energy efficiency**
  - **Networks exhibit tremendous operational diversity (topology, source behavior, desired quality, environmental conditions, instantaneous state)**
- **Objective: Data gathering energy should scale with desired quality, environmental conditions and internal state**
- **Scenarios: <Environmental Noise, Energy Vector>**
- **Point systems: All <Noise, State> *protocols***

Environmental Awareness

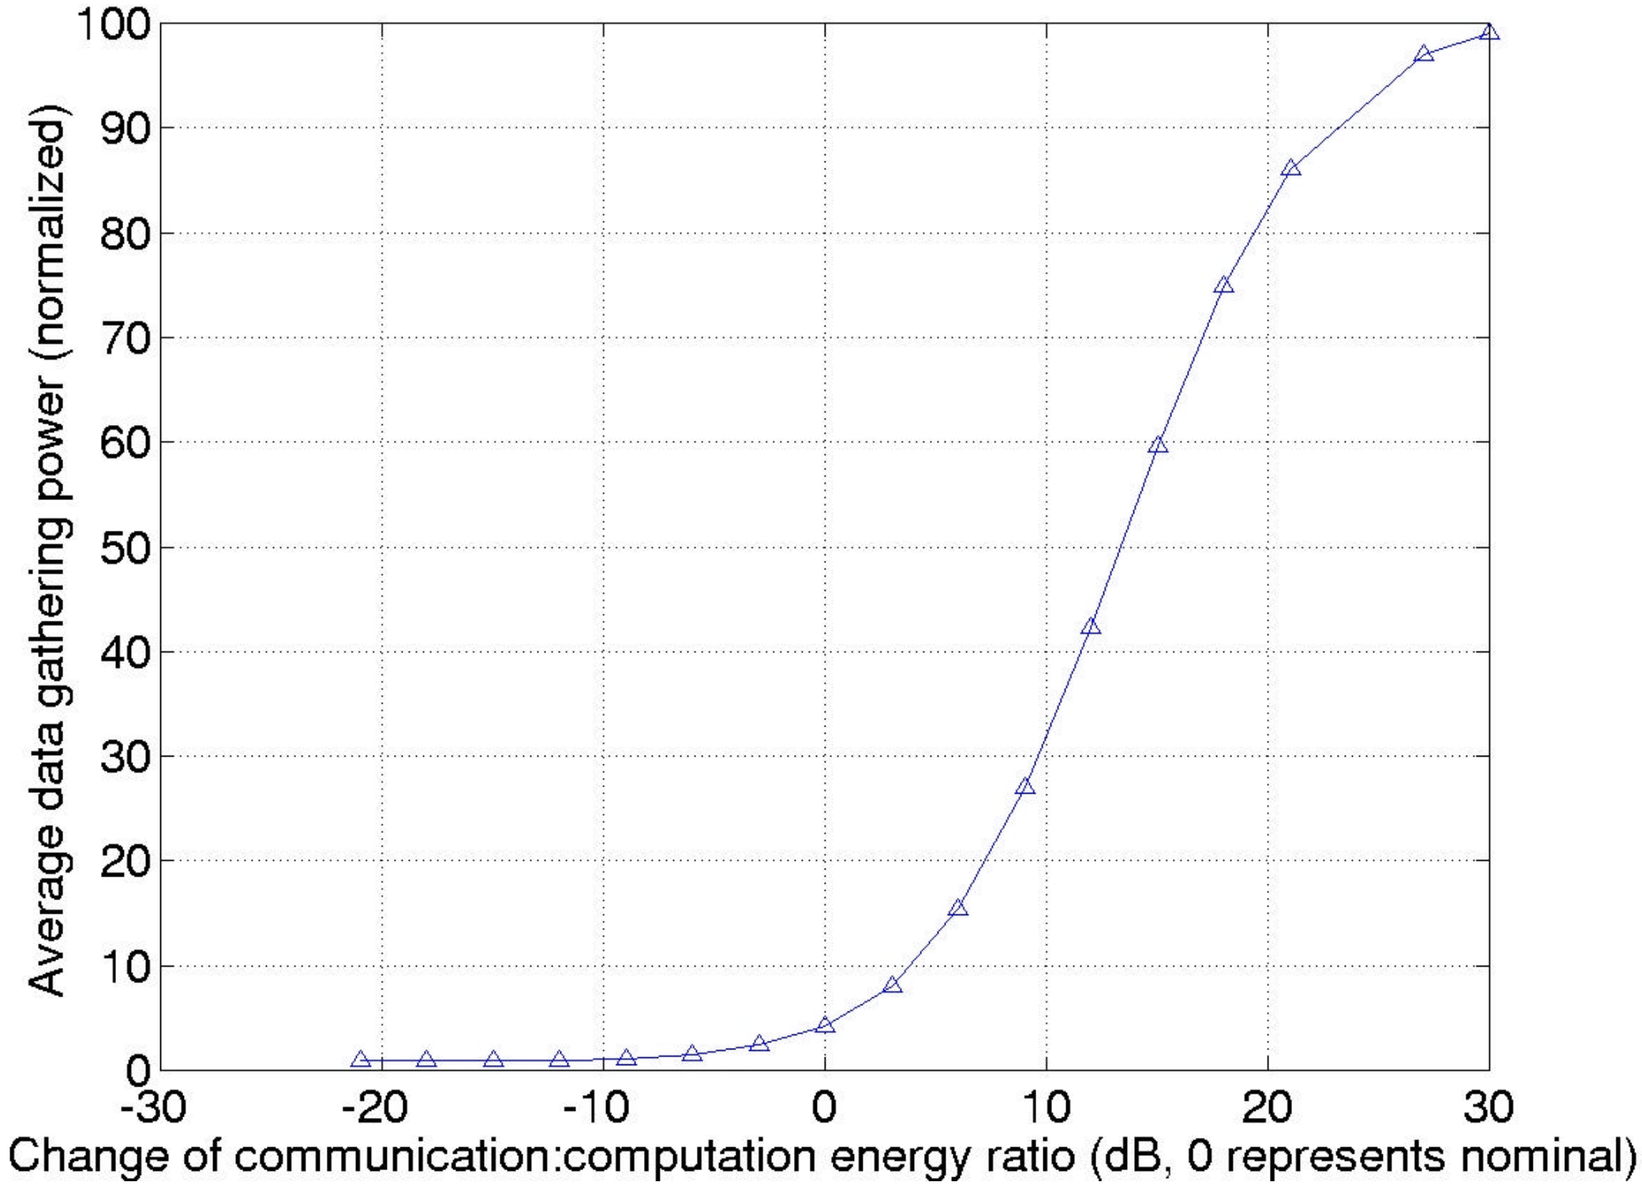

Protocol is potentially 10x more power-aware!

Awareness to State

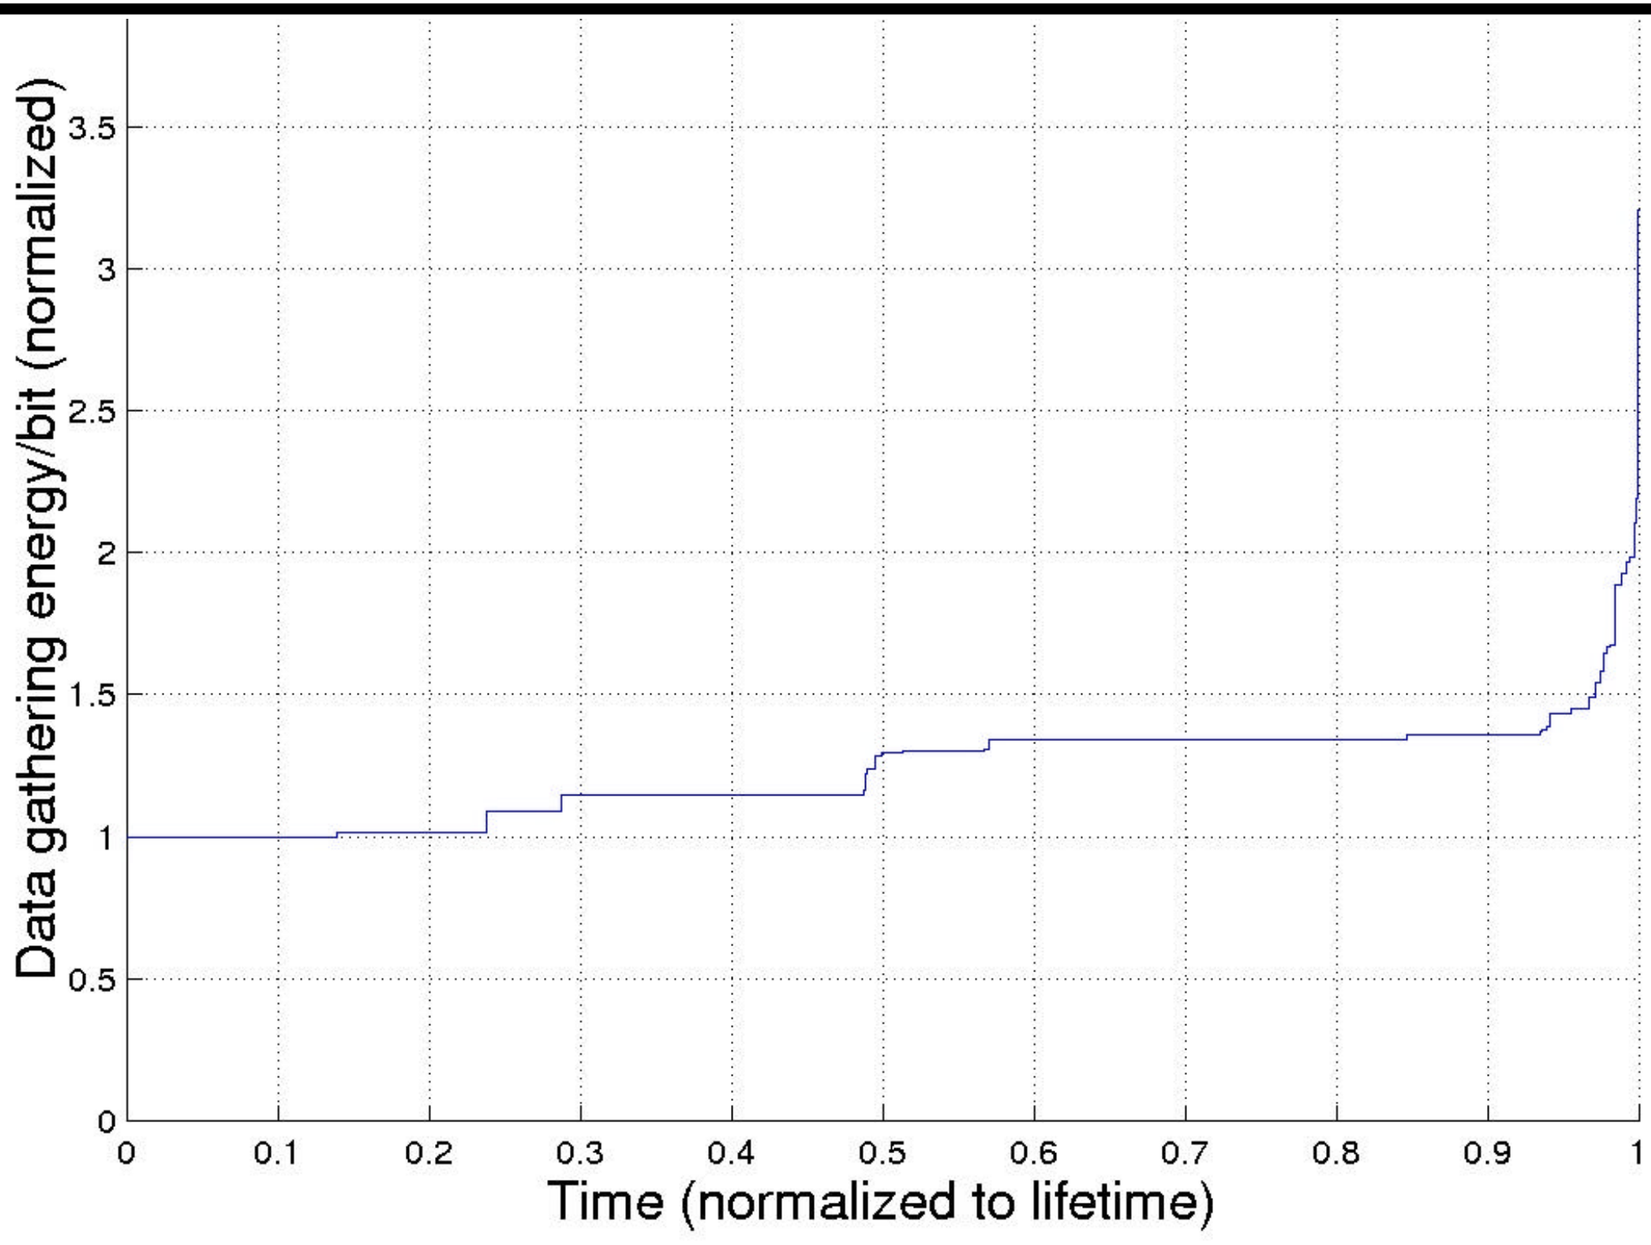

Protocol 2x more power aware than unaware versions

Summary

- Power-aware design can significantly enhance lifetime of battery constrained systems
- Power-awareness is a system-wide design philosophy
- Systematic methodology for power-aware design:
  - Characterize scenarios by understanding the awareness dimensions of a domain
  - Gather statistics and construct scenario distributions
  - Construct optimal ensembles
  - Measure power-awareness
  - Iterate
- Power-aware design is NOT low-power design
  - Low power design focuses on engineering point systems
  - Power-aware design focuses on characterizing and harnessing diversity by actively adapting the system