Exploring Tradeoffs in Accuracy, Energy and Latency of SIFT in Wireless Camera Networks

Teresa Ko, Zainul M. Charbiwala, Shaun Ahmadian, Mohammad Rahimi, Mani B. Srivastava, Stefano Soatto, Deborah Estrin
Center of Embedded Networked Sensing, Vision Lab, Networked Embedded Systems Laboratory, UCLA

Problem Description: Determine design tradeoffs for vision-based sensing systems.

Context

- Sensor systems collect & interpret sensor data.
- Intuitively, interpreting sensor data locally is better than at the server.
  - Scalability
  - Lower latency
  - Lower energy
- Image capture and transmission are on the same magnitude of energy consumption.
- The complexity of interpreting images relative to transmission is unknown, and is dependent on the application.

Conclusion

- Generic SIFT is not more efficient in terms of energy and latency than transmitting an image to a server and processing there.
- Application knowledge can result in changes in the location of computation and type of computation for more optimal behavior.
- Image processing/transmission dominates energy and latency budget.

System Model

System Block Diagram

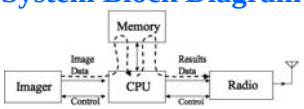

System Variables

- **Architecture**
  - Arithmetic Precision: Floating Point vs 16-bit Fixed Point
  - CPU Frequency: 50 MHz → 600 MHz
- **Application**
  - # of octaves: all, N-1, single
  - Scale space sampling: direct, indirect

System Operation and Data Flow

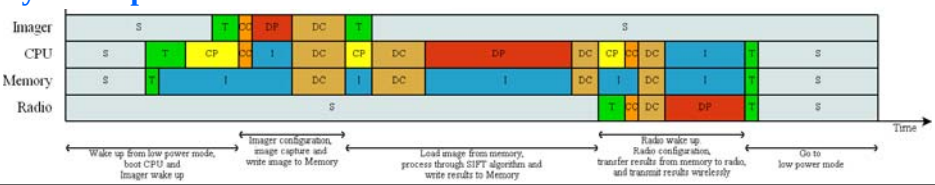

State-wise Power Consumption

Power (mW) / State	Imager	CPU 50MHz	CPU 600MHz	Memory	Radio
Sleep	1	0.081	0.081	0.018	0.054
Transition	50	13.2	141.3	171	48.8
Data Proc	42	20.8	264	0	47
Data Comm	42	39.8	283	171	48.8
Control Proc	10	20.8	264	0	50.76
Control Comm	42	22.9	266.1	0	48.85
Idle	10	11.2	37.2	0.360	50.76

Experimentation Results

Experimentation

- Ported SIFT to Analog Devices Blackfin DSP (BF-533) w/ floating point & fixed point arithmetic precision
- Built a system model for camera sensor node.
- Evaluated accuracy from experiments on PC using real life data set.
- Obtained cycle counts using Blackfin simulator
- Used the system model to extract energy consumption & latency from cycle counts.
- Interpreted tradeoffs between accuracy, energy, and latency.

Computational Breakdown

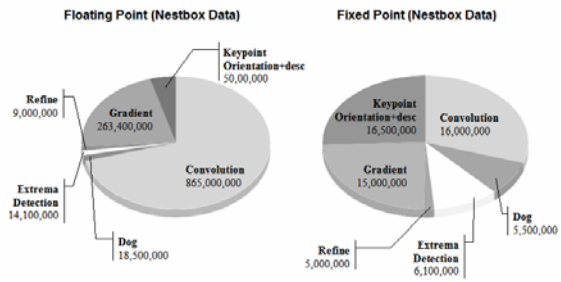

Energy and latency when varying arithmetic precision and # of octaves

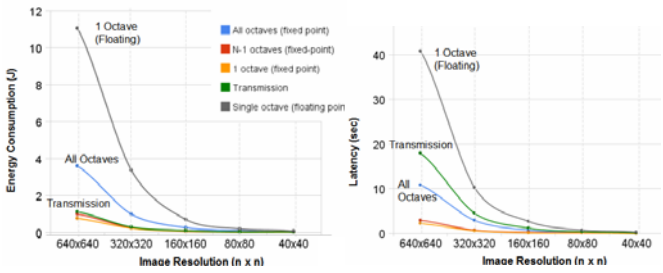

Effects on accuracy when varying arithmetic precision

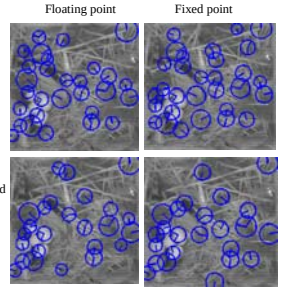

Difference in SIFT Descriptors

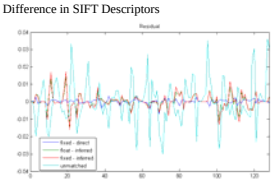

	Direct	Inferred	
	Fixed	Float	Fixed
% Matches	76.29%	64.72%	58.85%
% Misses	23.71%	35.28%	41.15%
% Extras	17.52%	41.51%	37.51%
Ave. position error	4.3306	12.4732	12.7448
Ave. orientation error	2.8680°	3.4028°	4.0723°
Ave. descriptor error	0.0407	0.0501	0.0528

Energy and latency when varying CPU frequency

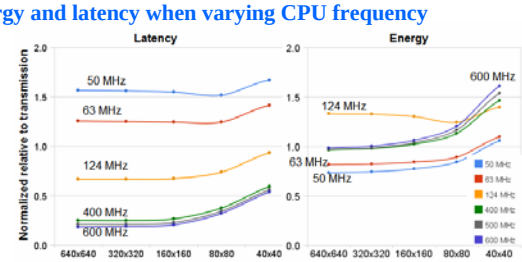

SVM Classification Results

	Float Direct	Fixed Direct	Float Inferred	Fixed Inferred
Precision	31.43%	29.23%	24.81%	24.77%
Recall	46.15%	43.51%	51.59%	53.54%
Distance	1.25	1.26	1.31	1.36