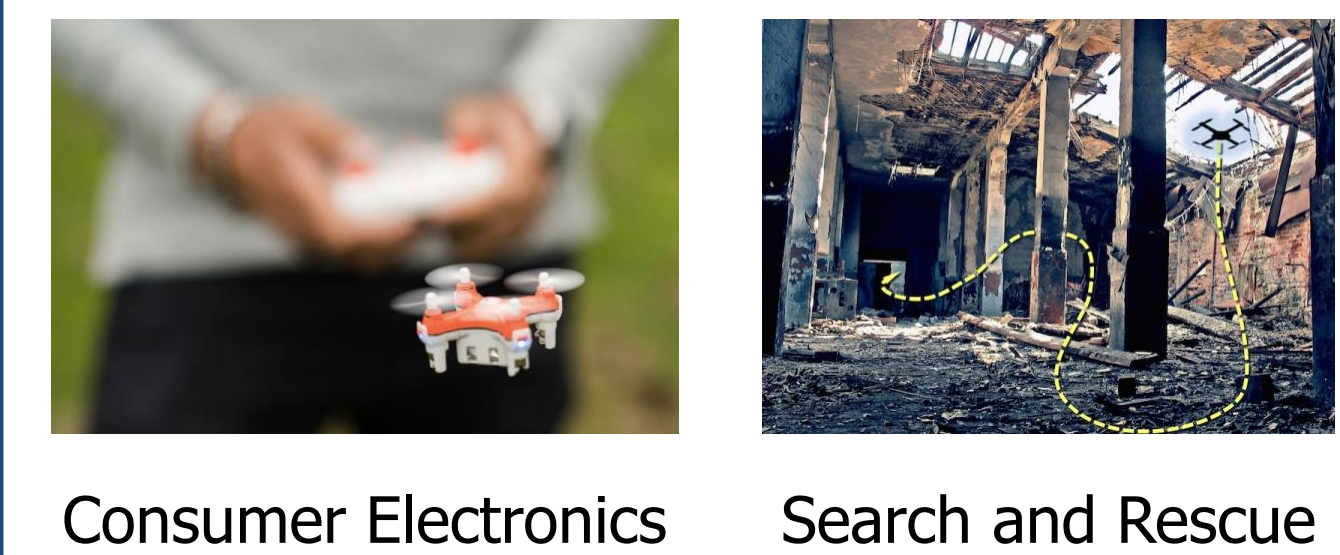



Visual-Inertial Odometry on Chip: An Algorithm-and-Hardware Co-design Approach

Zhengdong Zhang*, Amr Suleiman*, Luca Carlone, Vivienne Sze, Sertac Karaman
Massachusetts Institute of Technology

Motivation



Fully-autonomous navigation without a map is essential to these applications



Goal: Running fully-autonomous navigation without a map LOCALLY

Challenge: **Power and Speed**

	Lifting	1 W
	Cameras	1 W
Bottle-cap-sized nano UAV		
Goal	Desktop CPU	Embedded CPU
Keyframe rate	> 5 fps	8.4 fps
Power	< 2 W	28.2 W
		2.5 W

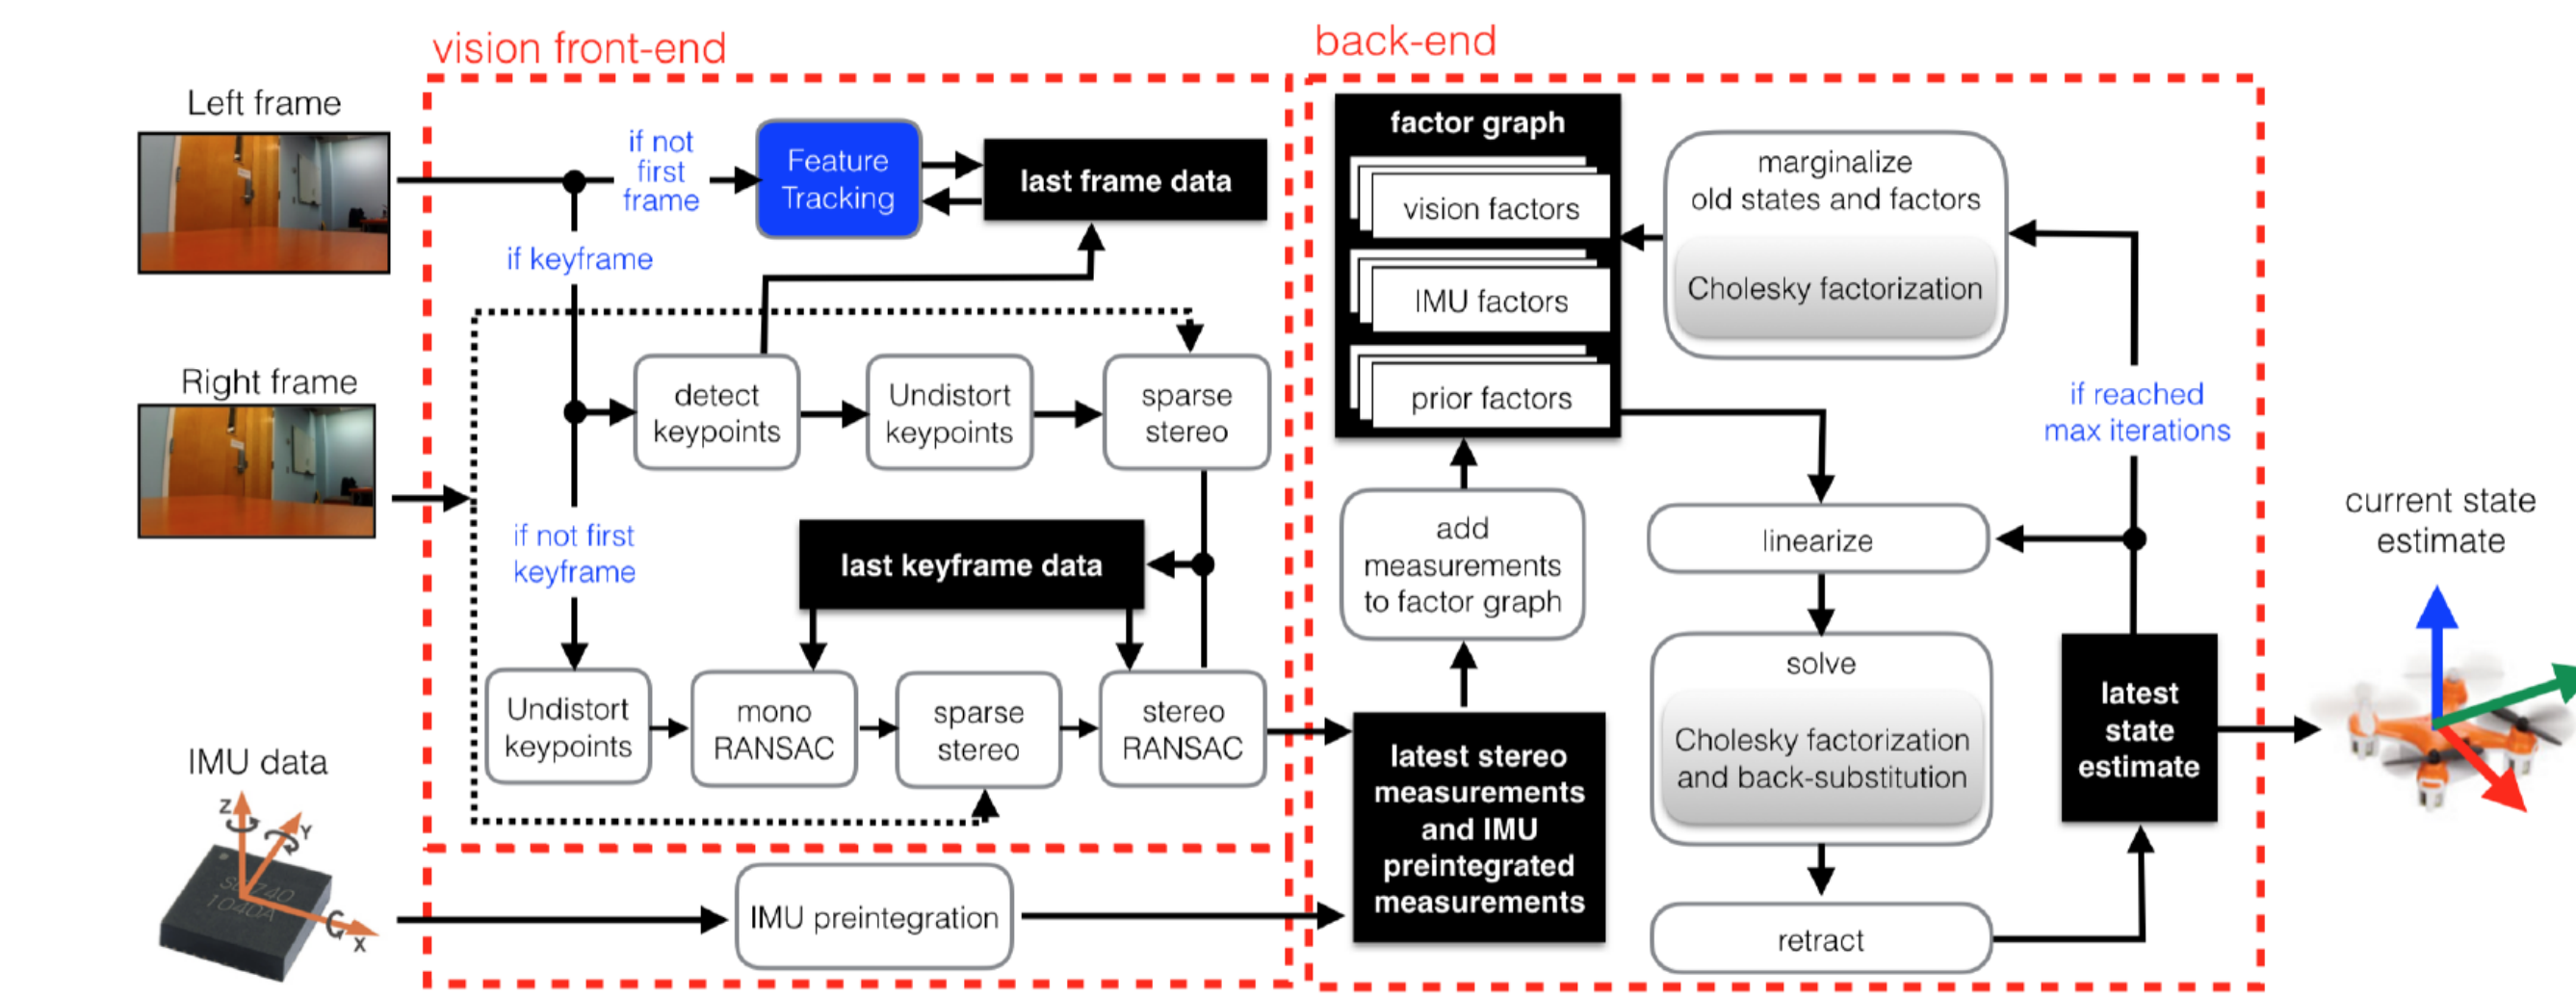
General Purpose Computing not good enough!

Low-power Specialized Hardware

	FPGA	ASIC
On-chip compute	1x (Reference)	6x
On-chip Memory	kB - MB	200x
Off-chip Memory	DRAM (GB)	

Standard VIO algorithms do not fit, we need an algorithm-and-hardware co-design approach

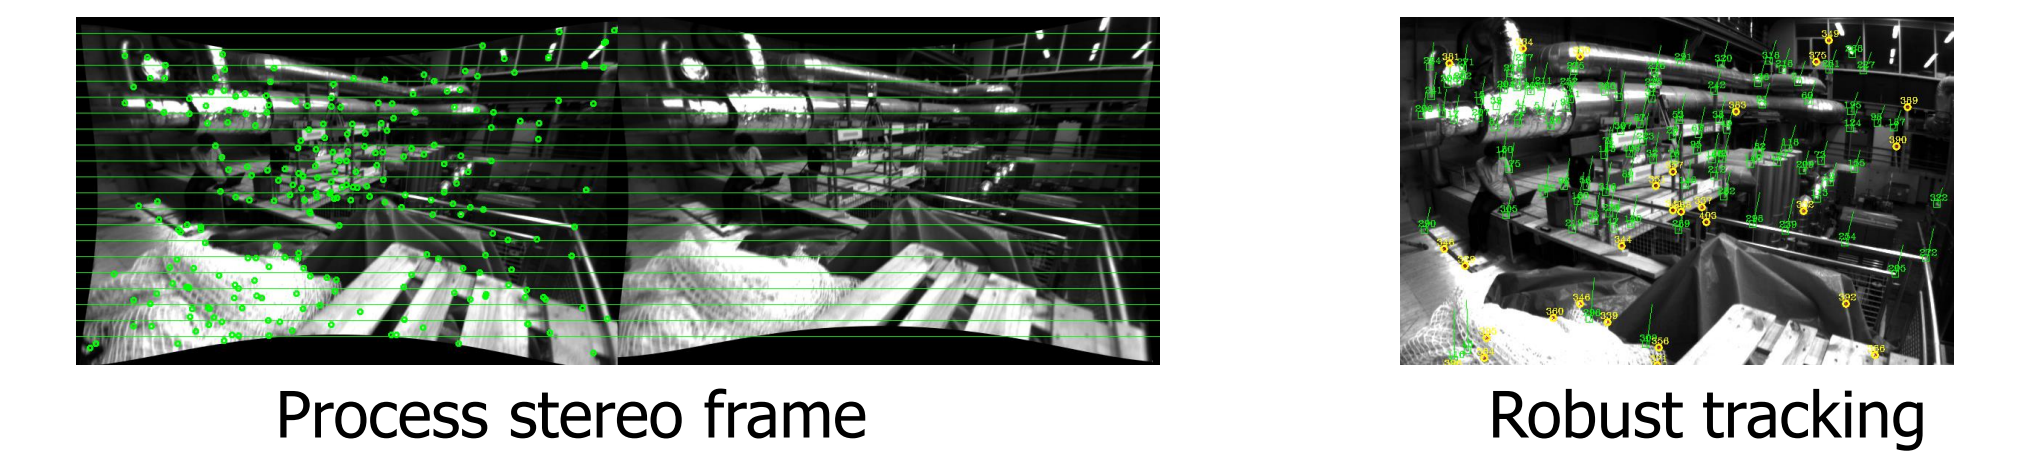
Visual Inertial Odometry (VIO) Algorithm



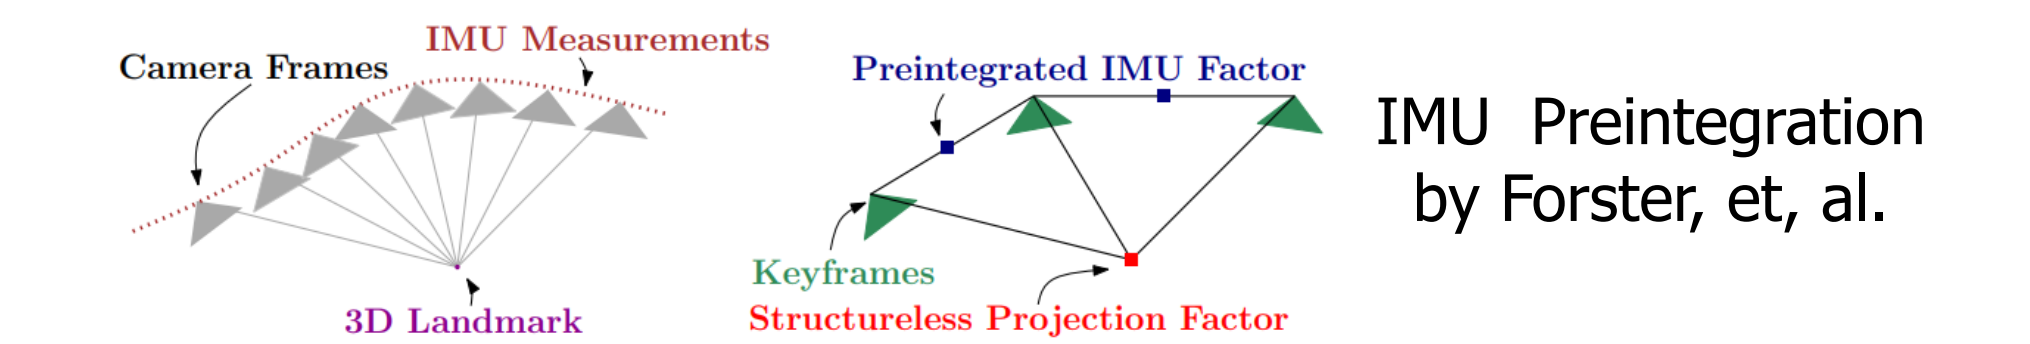
$$\min_x \sum_{(i,j) \in \mathcal{F}} \|r_{\text{IMU}}(x, \Delta R_{ij}, \Delta \tilde{p}_{ij}, \Delta \tilde{v}_{ij})\|^2 + \sum_{k \in \mathcal{L}} \sum_{l \in \mathcal{P}_k} \|r_{\text{CAM}}(x, l_k, u_{ik}^l, u_{lk}^r)\|^2 + \|r_{\text{PRIOR}}(x)\|^2$$

$x_i = (R_i, p_i, v_i, b_i)$: state consisting of robot poses, velocities and IMU bias
 $x = \{x_{t-h}, x_{t-h+1}, \dots, x_t\}$: extended state within the smoothing horizon
 $r_{\text{IMU}}(x, \Delta R_{ij}, \Delta \tilde{p}_{ij}, \Delta \tilde{v}_{ij})$: negative log-likelihood of the IMU measurements
 $r_{\text{CAM}}(x, l_k, u_{ik}^l, u_{lk}^r)$: negative log-likelihood of the vision measurements

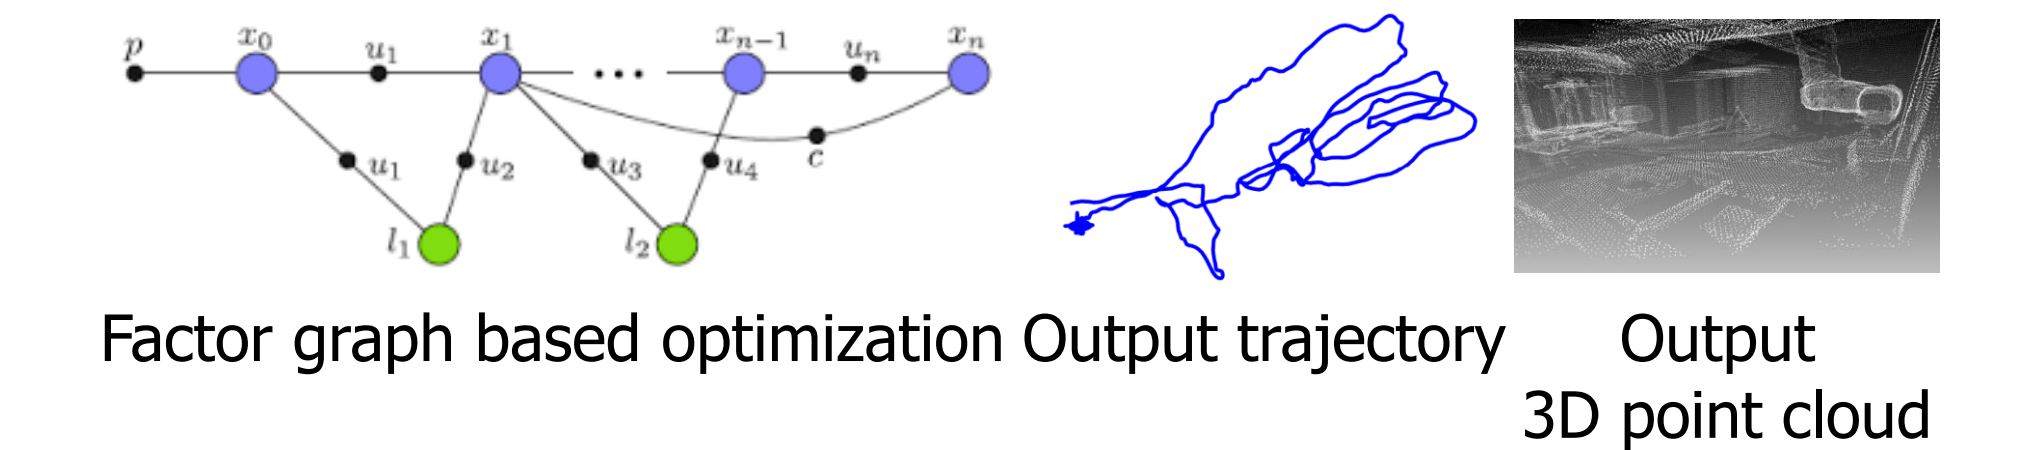
Vision Frontend



IMU Frontend



Backend



Algorithm-Hardware Co-design

Step 1: Specify Performance and Resource Goals

Battery life, endurance	Form factor	Accuracy	Speed, agility
Power $\leq 2W$	Board size, weight	Error ≤ 25 cm	Keyframe rate ≥ 5 fps

Step 2: Define Design Space, D

$D = H \times A \times I \times P$

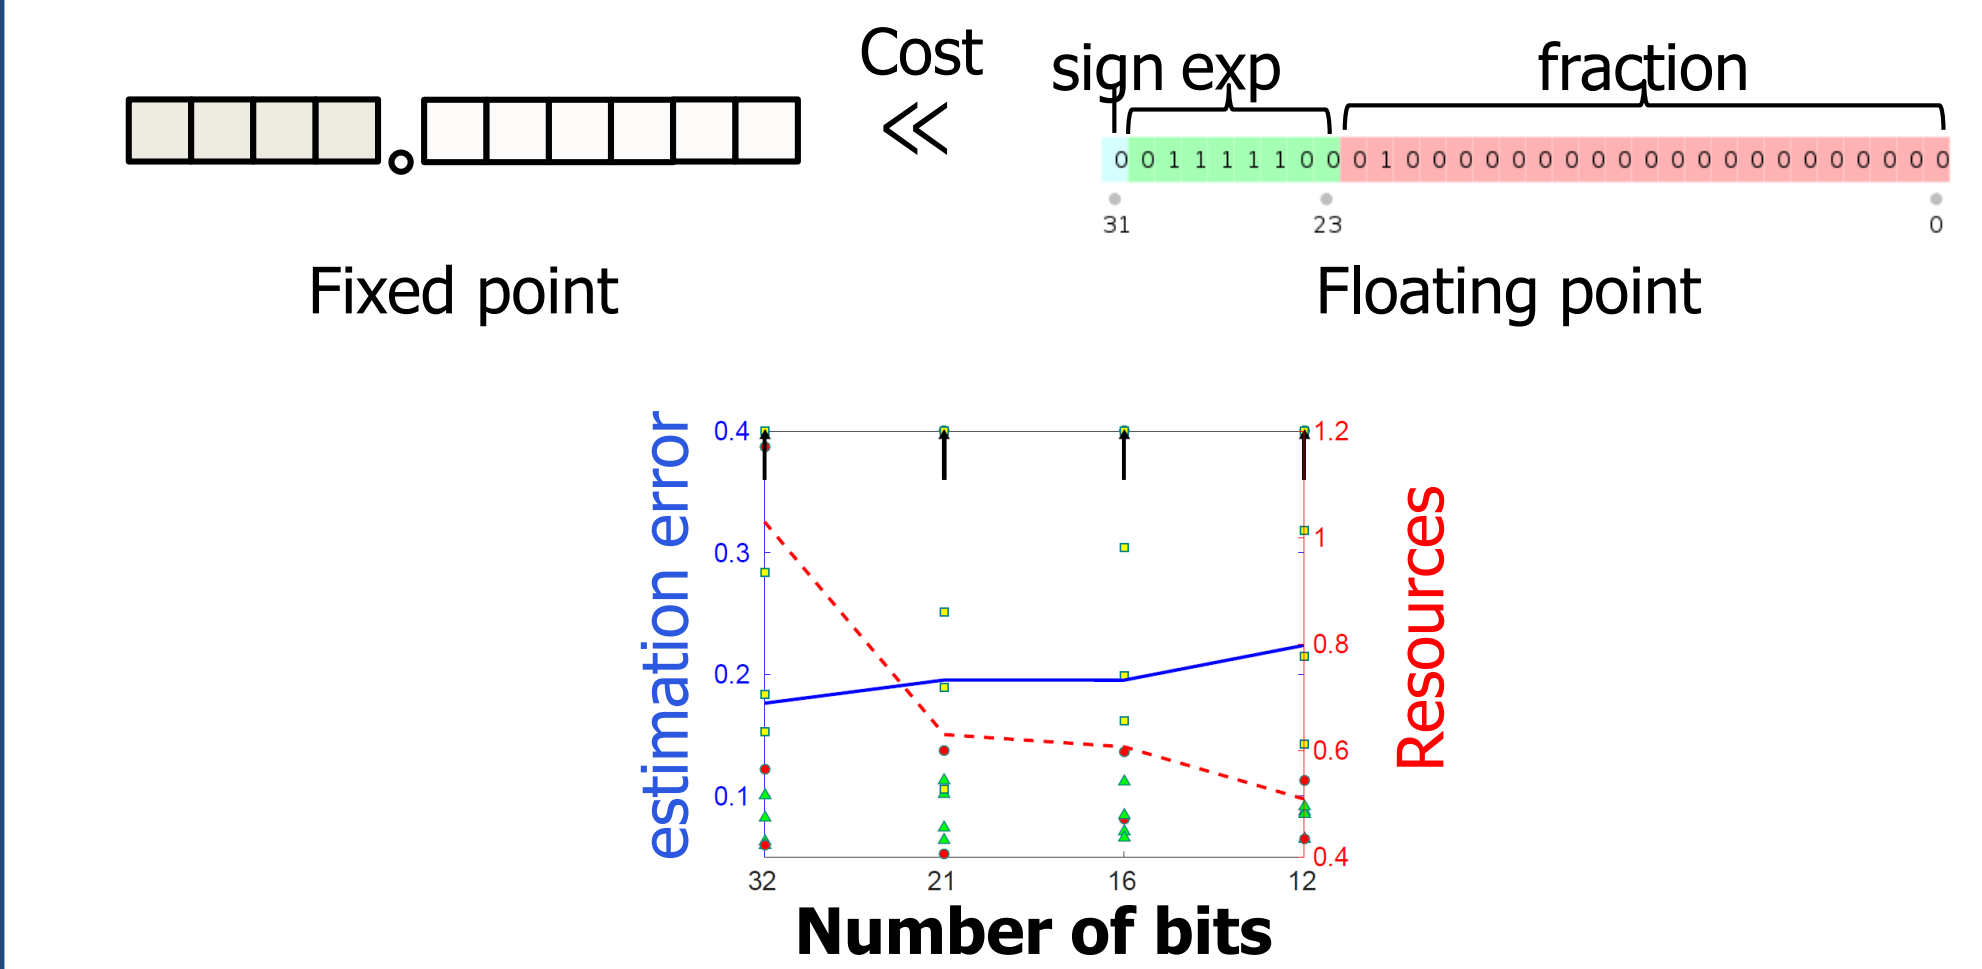
H	A	I	P
Hardware choices	Algorithm choices	Implementation choices	Parameter choices
desktop-CPU embedded-CPU FPGA ASIC	Tracking RANSAC Relinearization for Marginalization	Streaming Pipelining Parallelism Reduced precision	Max feature num Template size Max tracking levels Intra-keyframe time Nr. GN iterations

Step 3: Explore Design Space

- Choose a hardware (H):
- Split $A \times I \times P$ into (A, P_a) and (I, P_h)
 - Choose an algorithm a , and algorithmic parameters p_a in (A, P_a) to minimize power, while preserving the error
 - Choose an hardware implementation i , and hardware parameters p_h in (I, P_h) to re-establish the speed
- Iterate until finding a feasible design
Iterative Splitting Co-design

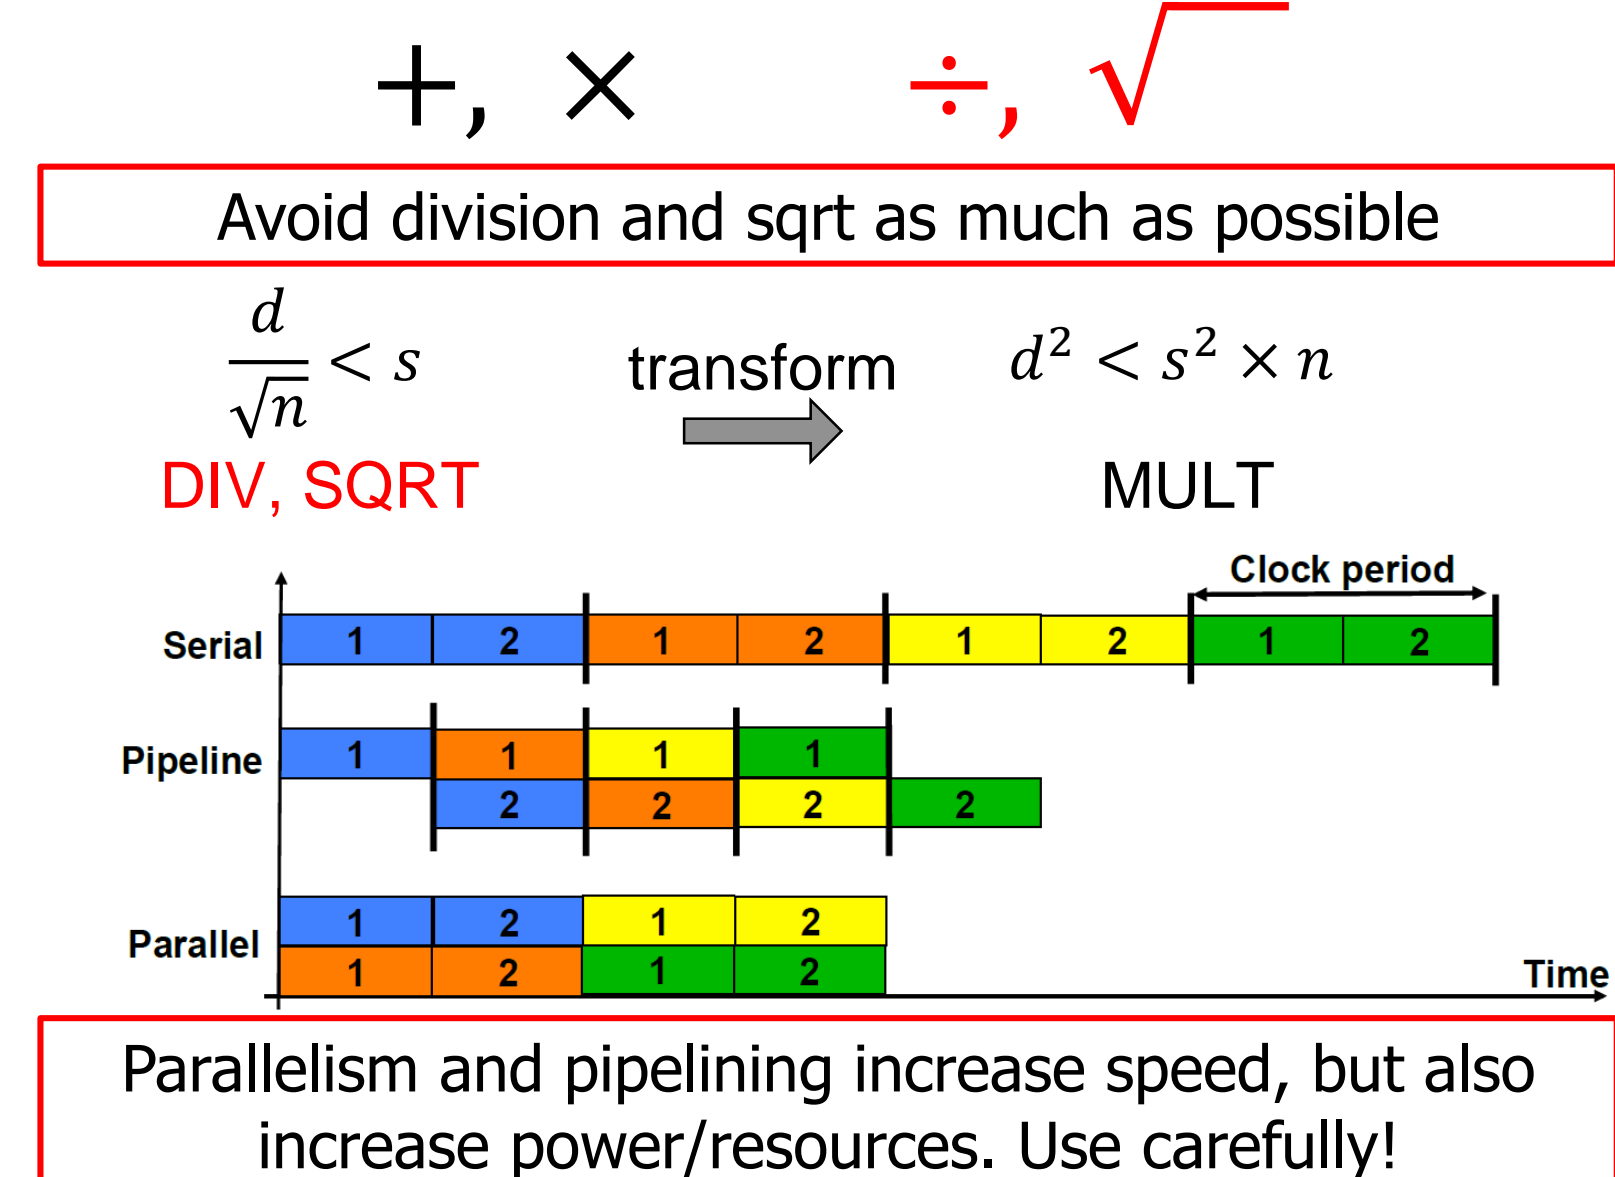
Co-designing VIO System on Chip

Reduced Precision of Data Representation



Reduce vision front-end to 16 bits fixed-point for efficient accuracy vs. memory trade-off

Hardware Design Choices




5-point RANSAC vs 2-point RANSAC

The vision front-end uses mono-RANSAC to remove outliers in the tracking

5-point RANSAC	2-point RANSAC (ours)
<ul style="list-style-type: none"> Complicated Hardware unfriendly 145 iterations 	<ul style="list-style-type: none"> Use the rotation from the IMU preintegration Hardware friendly 16 iterations

Co-design Results

Platform	Design Goal		Baseline		Design (a, p)		Design (h, a, i, p)	
	Accuracy	Power (W)	Desktop CPU	Embedded CPU	Desktop CPU	Embedded CPU	FPGA	
Accuracy	≤ 0.25			0.15		0.16		0.19
Front-end Throughput (fps)	≥ 20		15.4	3.9	20.8	5.2		20
Back-end Throughput (fps)	≥ 5		8.4	2.0	12.7	2.7		5
Power (W)	≤ 2		28.2	2.5	26.1	2.3		1.5



Resource utilization on Xilinx Kintex-7 XC7K355T FPGA: Only **2.1 MB** memory is needed

Resource	Front-end	Linear Solver	Linearize	Marginalize	Stereo Factors	IMU & Other Factors	State Estimation	Total	Utilization
Memory	1.5 MB	355 KB	0	180 KB	86 KB	15 KB	3 KB	2.1 MB	67.3 %
Block RAM	378	80	0	41	22	4	1	0.5 K	73.5 %
Flip Flops	72 K	5 K	67 K	Shared	n/a	n/a	n/a	144 K	32.4 %
LUTs	111 K	16 K	65 K	Shared	n/a	n/a	n/a	192 K	86.2 %
DSP	607	62	102	Shared	n/a	n/a	n/a	771	53.5 %