Stereo and IMU Assisted Visual Odometry on an OMAP3530 for Small Robots

Steve Goldberg and Larry Matthies
Computer Vision Group
Jet Propulsion Laboratory
California Institute of Technology

indeliblesteve@gmail.com
lhm@jpl.nasa.gov

What We Want to Enable

- 3-D perception, localization, and other functions, fused with IMU, on small, unmanned ground and air vehicles for recon, EOD, and other missions

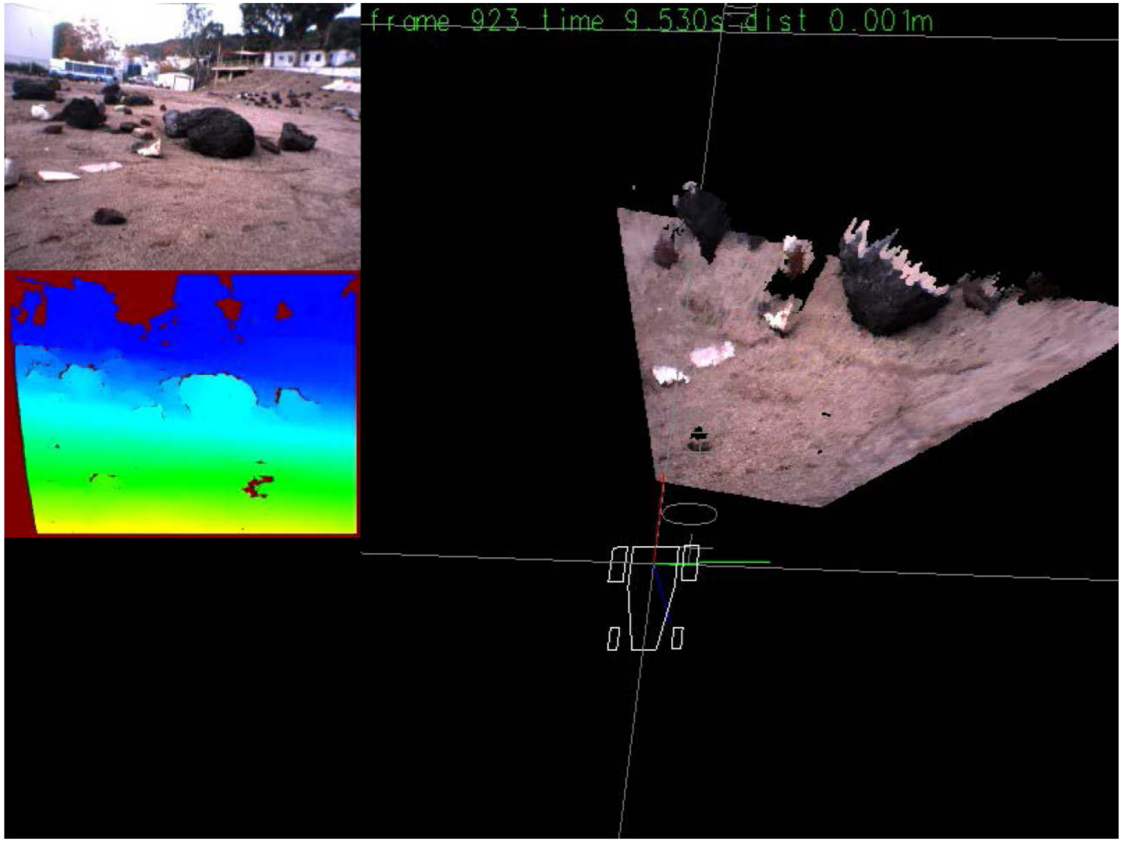

On What Types of Vehicles?

- Ground and air at a range of scales from > 7 kg to < 50 g

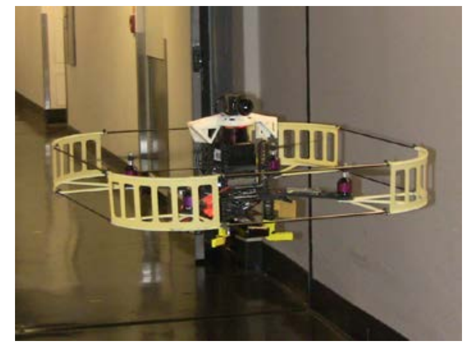

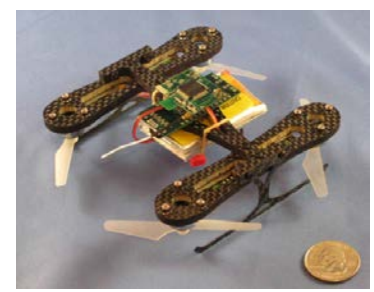

Processor Options

Jet Propulsion Laboratory

- Intel Atom: too big and hot
- GPU's: too big and hot
- ASIC's: too expensive, long development, inflexible
- FPGA's: relatively long time to program
- DSP's: good for some things
- Smart phone SoC's, e.g. TI OMAP family
  - DSP plus ARM
  - Already in popular board (Gumstix) used in robotics
  - Already in some research and commercial robots (e.g. FirstLook)
  - Good roadmap

OMAP and DaVinci Processor Family

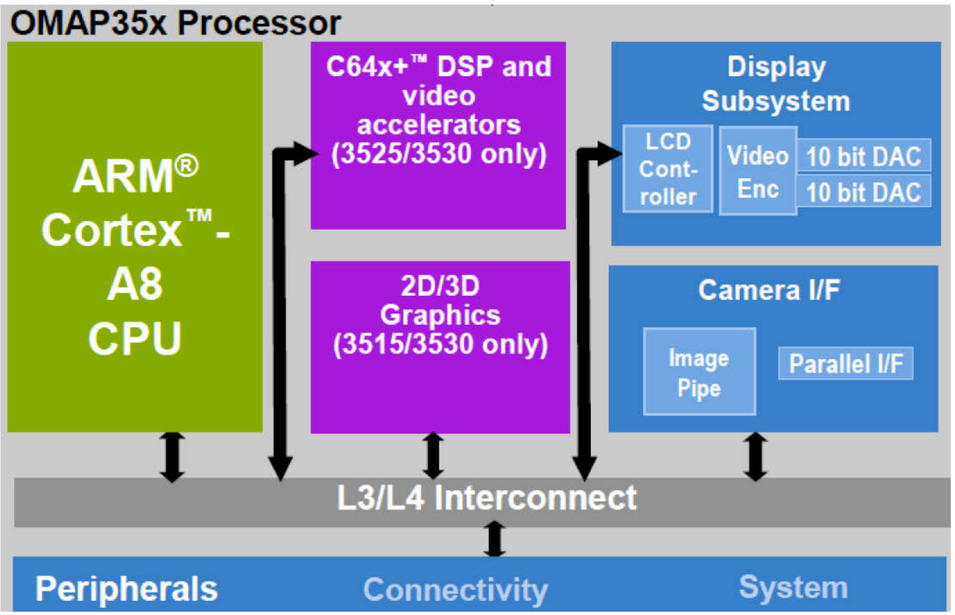

Gumstix Overo Fire
OMAP3530
6 cm, 6 g, < 3 W

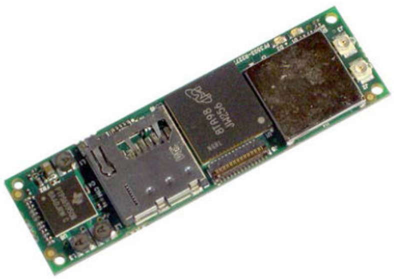

LogicPD Torpedo
OMAP3530 or DM3730
3 cm, 2 g

- Up to 1 GHz ARM, 800 MHz DSP
- OMAP4: two ARM Cortex-A9 cores, up to 1.5 Hz
- OMAP5: two ARM Cortex-A15 cores, up to 2 GHz; two ARM-M4 cores; up to four simultaneous camera inputs
- Related DaVinci chips with floating point (C67) instead in integer DSP
- LANdroid robot has OMAP3530; FirstLook robot has DM3730

Objectives

- Integrate two cameras and an IMU with an OMAP3530 and ensure very accurate synchronization and time-stamping of all sensors
- Implement dense stereo vision and visual odometry on the OMAP3530 to establish what speed is possible
  - Visual odometry on ARM side to exploit floating point
  - Stereo on DSP side
- Exploit IMU to aid feature tracking and state estimation

Processor Board: Gumstix Overo Fire

- Pros:
  - OMAP3530 support
    - 720 MHz ARM
    - 520 MHz C64x+ DSP
  - Small and light:
    - 17 x 58 x 4.2 mm @ 5.3 g
  - MicroSD card built in
  - WiFi optional
  - Peripheral boards
    - USB, Ethernet, HDMI
- Cons:
  - Rumored flaky connectors
  - Slow MicroSD
  - 1.8v GPIO
  - One camera port

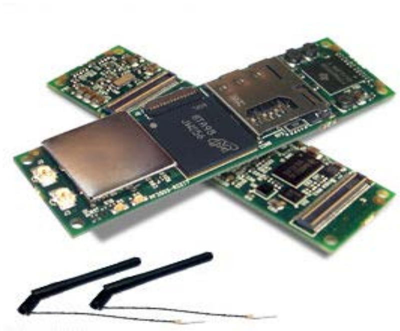

Cameras: Point Grey FireFly MV USB

- Global shutter CMOS
- DC1394 support over USB
- Synchronized via onboard GPIO
- 720x480 @ 60 fps
- Format7
  - ROI, 2x2, 2x1
- Equivalent CMOS imager in Gumstix Caspa cameras

- Pros:
  - Easily integrated
  - Global shutter
- Cons:
  - Wastes power
  - USB buggy

Inertial Sensor: CHR-6dm AHRS

- 3 axis rate gyros
- 3 axis accels
- 3 axis magnetometer
- ARM Cortex-A3
  - Floating point
  - 76 MHz
- Low power
  - 160-260 mW
- OpenSource firmware

- Pros:
  - Inexpensive: $200
  - OpenSource
  - GPIO
- Cons:
  - Low grade sensors
  - Euler EKF unstable
  - Requires Windows

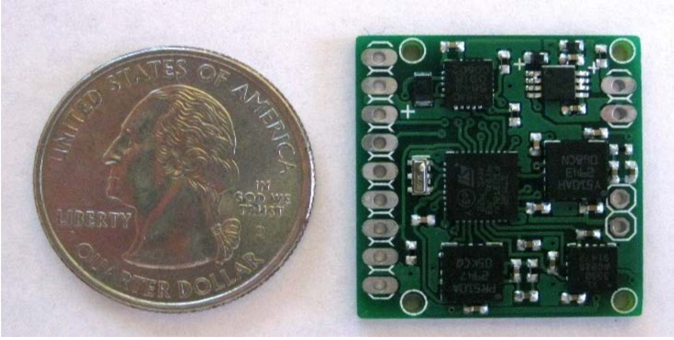

System Description

- Stereo and VO
- 512x384x51 @ 14 fps
- 320x240x32 @ 30 fps
- 200 features selected
- Good feature tracking
- Real-time display
- Optional WiFi, Ethernet, USBNet
- Tightly integrated
  - µsec camera to IMU sensor timing
- ARM Linux
- Upgradeable (DM3730)

Algorithms Implemented: SAD1 Dense Stereo

- Ref. Goldberg, Maimone, Matthies, IEEE Aerospace Conf, 2002
- SAD1 local optimization correlator with rectification, difference of boxes prefilter, LR check, subpixel estimation

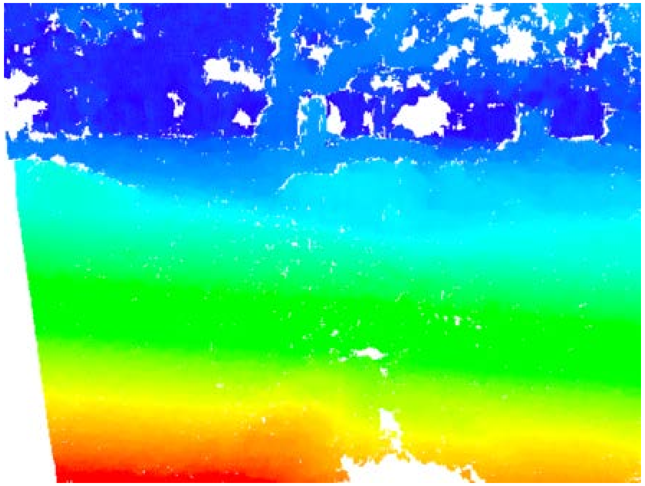

Algorithms Implemented: Visual Odometry

- Ref. A. Howard, IROS 2008
- Built on top of real-time dense stereo algorithm
- FAST feature detection to extract about 200 features
- Discard features without range data from stereo depth map
- Feature tracking (from Howard):
  - Detect FAST features in next frame, associate candidates, select best candidate with correlation scores
- Egomotion estimation:
  - Rigidity test to reject outliers, then minimize reprojection error with data editing loop
- IMU-assisted tracking:
  - Use gyros to predict feature locations and do local search

Illustration of Visual Odometry Performance

- VO tests in JPL arroyo
- Mixed open/vegetated terrain
- Multiple loops; ~100m/loop
- Error < 1.5m over 300m
- OMAP implementation validated with consecutive pairs of such images

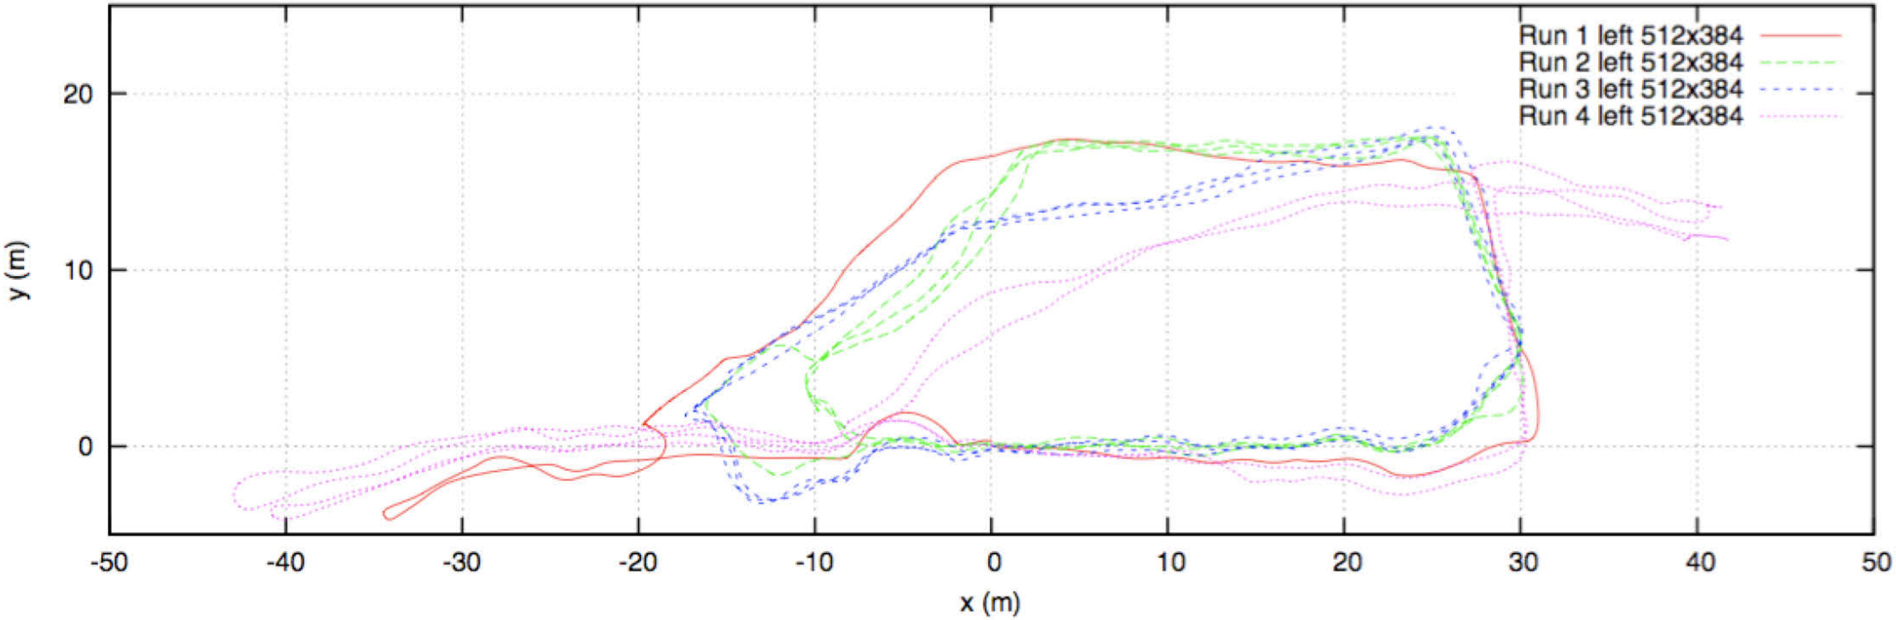

Software Architecture

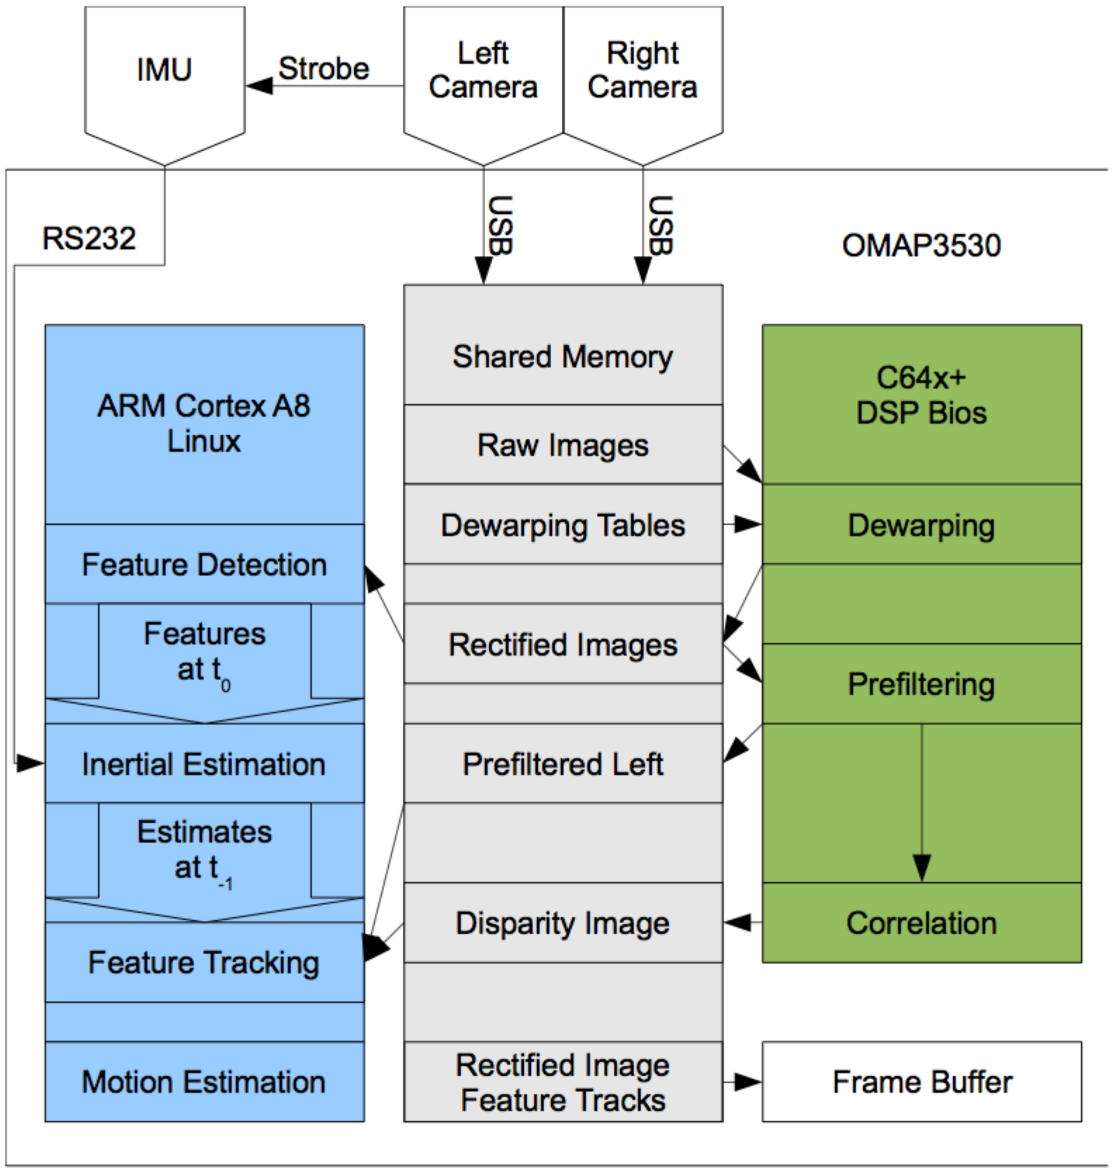

Results: Stereo and Feature Tracking (512x384)

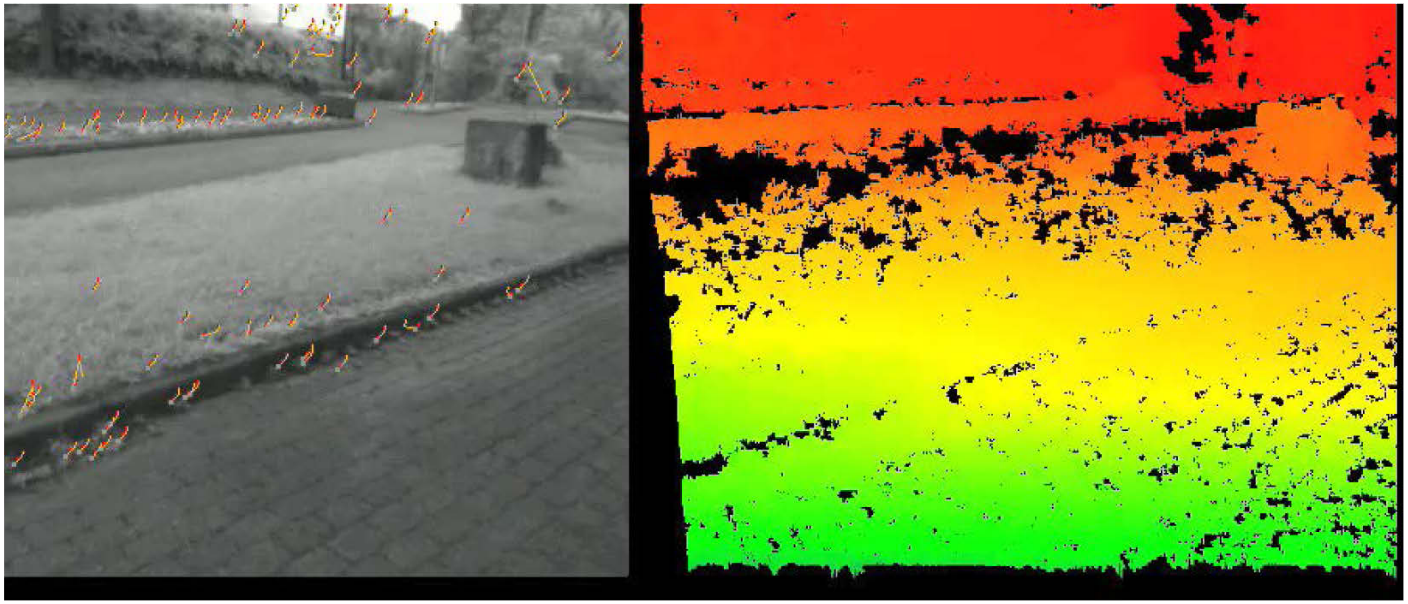

Related Work

- Stereo:
  - Khaleghi, ECVW 2008; 3x3 Census on Blackfin DSP
  - Humenberger, CVIU 2010; 8x8 Sparse Census on C64
  - Ambrosch, ICEEEI 2010, Sparse Census on C64's
  - Chang, ICME 2007; 4x5 Jigsaw SAD on C64 DSP
- Visual odometry:
  - Many algorithms done in software
  - Not much (anything?) published on low power embedded versions
- IMU-assisted feature tracking:
  - Exemplified by Hwangbo, IROS 2009; IMU used to predict feature locations and pre-warp image templates in mono imagery for UAV application

Runtime: Stereo Depth Map

Stereo	320x240x32	512x384x51	640x480x60
Overhead	3ms	6ms	10ms
Rectification	0.5ms	2ms	3ms
Box-filter	1ms	3ms	5ms
Correlation	14ms	55ms	98ms
Sub-pixel	3ms	5ms	8ms
Frame Rate	46fps	14fps	8fps

Table 1: Stereo runtime performance on 520Mhz C64x+ side of OMAP3530

DSP Stereo	Method	Mde/Ghz	Mde/J
Khaleghi [27]	3x3 Census	19.2	5
Humenberger [17]	8x8 Census	78.38	15.68
Ambrosch [12]	Sparse Census	120.23	n/a
Chang [30]	4x5 Jigsaw SAD	88.47	n/a
OMAP3530	7x7 SAD	280.77	58.8
DM3730	7x7 SAD	264.96	84.8

Table 2: Comparison of DSP stereo runtimes and power utilization.

Runtime: Feature Tracking

Jet Propulsion Laboratory

Initial Features	200	300	400
Detection	7ms	7ms	8ms
Matching	6ms	12ms	23ms
Inlier	13ms	24ms	50ms
Motion Estimation	16ms	20ms	21ms
Inlier Features	80	115	140
Frame Rate	14fps	10fps	7.5fps

Table 3: 512x384 visual odometry runtime breakdown on 720Mhz ARM side of OMAP3530

Summary, Conclusions, Future Work

- Implemented dense, SAD1-based stereo vision and visual odometry with IMU-assisted feature tracking on OMAP3530
  - ARM does VO, DSP does stereo
- 46 fps stereo at QVGA or 8 fps at VGA resolution, while simultaneously tracking up to 200 features and doing visual odometry at up to 30 Hz
- Stereo implementation achieves 58.8 Mde/J, 3.75x faster than previous best reported
- Next steps:
  - Data logging capability nearly complete; then conduct further validation
  - Integrate on a UGV and test autonomous navigation with it
  - Adapt to micro UAV applications (monocular)