An Extended Kalman Filter for Cell Phone Orientation Tracking

Nisarg Kothari (ndk@andrew.cmu.edu)

***Abstract*—Modern smartphones are equipped with several sensors that can be used to track the orientation of the phone: accelerometers, magnetometers, and gyroscopes. This capability is useful for games, localization, and augmented reality applications. Achieving accurate tracking, however, is non-trivial due to sensor limitations. In this work, I present an Extended Kalman Filter to solve the orientation tracking problem on a smartphone. The filter tracks the orientation using quaternions to avoid the problems associated with Euler Angle singularities. The process model, measurement model, and noise model used for the filter are specified. The filter converges to and tracks the correct orientation in all tested scenarios.**

I. INTRODUCTION

Imagine a cell phone application that provided a 'window' into a virtual environment which changed realistically as the user moves the phone. Such an application is possible by tracking the orientation of the cell phone in space. Currently, many applications use low-pass filtered accelerometer and magnetometer values to perform orientation estimation. This method could be improved by employing a more sophisticated filter and by incorporating measurements from a gyroscope, a sensor which has only recently become available on cell phones.

This work discusses the implementation of an Extended Kalman Filter to track the orientation of a cell phone. An Extended Kalman Filter is a Bayesian method for nonlinear state estimation [4]. The orientation tracking problem in nonlinear due to the presence of rotations. Two important considerations for the filter are tracking the calibration offsets of the gyroscope sensor in the filter, and modifying the noise parameters for the accelerometer and magnetometer measurements depending on the deviation of properties of the signal from expected values.

II. COORDINATE FRAMES

There are two coordinate frames relevant to the Kalman filter: the global frame describing the orientation of the phone with respect to the Earth, and the local frame which describes quantities with respect to the phone itself. The accelerometer and magnetometer measure quantities that are defined in the global frame but reported in the local frame. The gyroscope measures rotation purely in the local frame. The purpose of the filter is to express orientation with respect to the global frame.

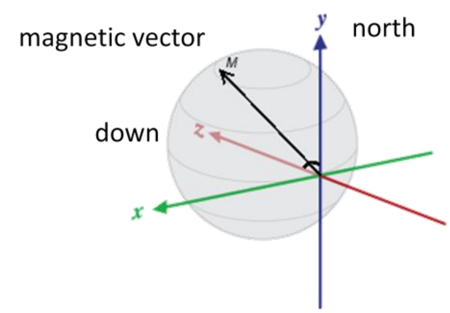

Fig. 1. Global (relative to earth) coordinate frame diagram.

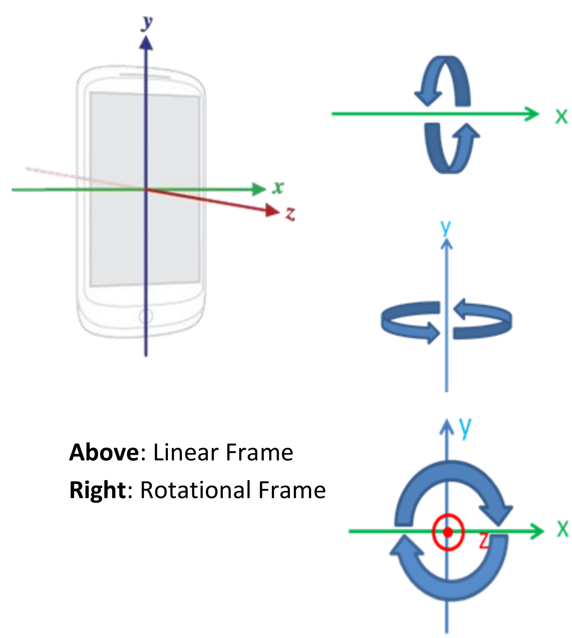

Fig 2.Local (relative to phone) coordinate frame diagram.

III. Kalman Filter Design

An early question in designing a Kalman Filter is what states it will track. Since I maintain orientation separately using a quaternion, the filter must track the angular velocity and the change in angle since the last quaternion update in its state. Some additional state is added to the filter due to the nature of the MEMS gyroscopes in cell phones. These gyroscopes have offsets in their reported values which cause them to report rotation even when they are completely still. Furthermore, the offsets change slowly over time, in particular due to changes in temperature. The filter is designed to track these offsets to minimize the amount of error introduced by a change in the calibrated offset. The complete state tracked by the filter is summarized below:

State tracked by Kalman Filter	Number of Values
Local Angular Velocity (x,y,z)	3
Change in local angle (Δx, Δy, Δz)	3
Gyroscope offsets (x,y,z)	3
Total Size of State	9

The nonlinearities inherent in tracking rotation are dealt with by numerically linearizing the process and measurement matrices as necessary, and by periodically updating the quaternion for changes in rotation using the small angle approximation and then zeroing those changes out from the Kalman filter state.

A. Process Model

Because the phone is likely to be actively manipulated by the user rather than passively evolving along a particular system trajectory, it is difficult to predict the future state with any degree of accuracy. I make very simple assumptions about the future behavior of the state and depend on the measurements to correct the prediction.

The first assumption is that the angular velocity of the phone will remain constant. Depending on the particular application the orientation tracking is used for, both a constant angular velocity and a zero angular velocity are potentially reasonable assumptions.

$$w'_{local} \leftarrow w_{local}$$

$$\Delta angle \leftarrow R w_{local} \Delta time$$

The second assumption is that the gyroscope offsets will remain constant. Since the offsets usually change slowly in reality, this is a reasonable assumption.

$$gyro\_offsets' \leftarrow gyro\_offsets$$

B. Noise Model

The design of the noise model is very important for an effective orientation tracking filter. The reason is that the accelerometer and magnetometer sensors regularly suffer from large disruptions that make the values they report useless for tracking orientation. The accelerometer can be disrupted by any linear accelerations exerted on the device, and the magnetometer can be disrupted by things like rebar, motors, electrical wire, and anything else which creates or interacts with magnetic fields.

Each sensor reading is checked for indications that it is suffering from disruption. If it is, the measurement noise for the reading is increased, and in extreme cases the reading is skipped altogether.

1) Acceleration Check

The magnitude of the acceleration is checked to see if it is in the vicinity of the expected 1g. The measurement noise of the particular measurement is increased in proportion to the deviation of the acceleration from 1g.

2) Magnetic Field Strength Check

The strength of the Earth's magnetic field for a given global location is roughly constant. Any significant deviations from this expected value must be due to magnetic interference. Similarly to the accelerometer reading, the noise of the magnetic field measurement is increased in proportion to the deviation of the magnetic field strength from the expected value, which is 55uT for the Northeastern United States [3].

3) Magnetic Inclination Check

In addition to an expected field strength for a given global location, there is also an expected inclination. The inclination of a magnetic field is the angle of the field vector with respect to the ground. The inclination varies from 90° at the poles to 0° at the equator. In the Northeast United States, it is roughly 70° [3]. The inclination I of the received magnetic field reading is calculated by using the current orientation estimate. If the orientation quaternion is expressed as [a b c d], with 'a' being the scalar value, then

$$Down = [2bd - 2ac, 2cd + 2ab, a^2 - b^2 - c^2 + d^2]$$

$$North = [a^2 + b^2 - c^2 - d^2, 2bc - 2ad, 2bd + 2ac]$$

$$\mathrm{I} = \mathrm{atan2}(North \cdot magnetic, Down \cdot magnetic)$$

where North and Down are unit vectors in those directions expressed in terms of local coordinates. They are calculated by extracting rows from the conversion of the orientation quaternion into a rotation matrix [2].

If the inclination deviates significantly from the expected value, the reading is rejected and is not entered into the filter.

C. Measurement Model

The accelerometer readings are predicted from the state in a straightforward manner using the orientation quaternion. The value reported by the accelerometer is assumed in the measurement model to be due only to the effect of gravity, with no linear accelerations present. Therefore, the reading is predicted to be the direction of 'down' expressed in the local coordinate frame.

I initially attempted to predict the complete magnetometer reading by assuming a constant value for the magnetic inclination (for a given global location). This lead to very large errors, and ultimately I decided to only attempt to predict the direction of the horizontal component of the magnetic field, the component pointing in the 'North' direction. To match this predicted value properly, the vertical component of the measured magnetic field is removed and the value is normalized before the error is calculated.

The gyroscope readings can be predicted with the most accuracy by the measurement model. The angular velocity of the device is assumed to be constant over time, in which case

$$gyro_{predicted} = w_{local} + gyro\_offsets$$

D. Kalman Filter Loop

This section describes the overall operation of the filter. There are seven different types of measurement steps corresponding to any combination of accelerometer, magnetometer, and gyroscope sensor readings. If a reading was not available, the corresponding rows were omitted when generating the measurement matrix. The filter was run at 60hz, and the individual sensors reported at roughly 30hz each.

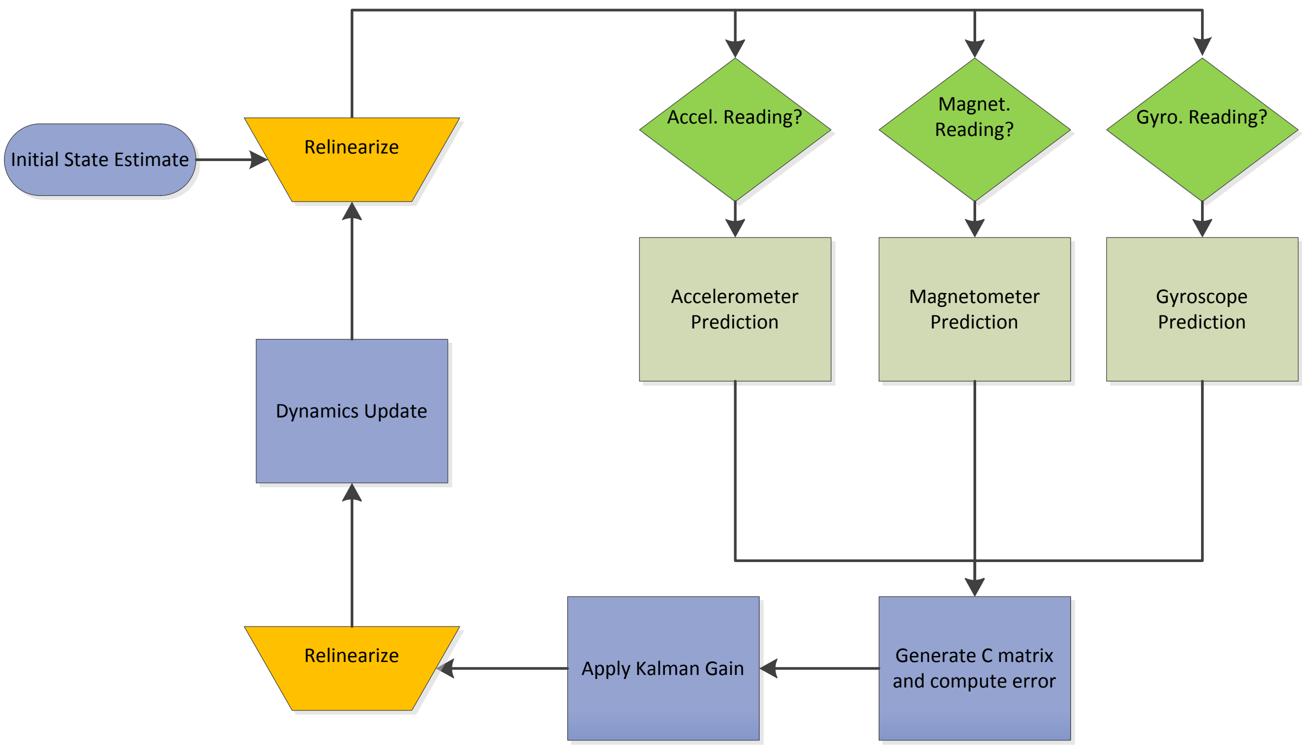

Fig 2. Diagram of Kalman Filter operation

IV. Experiments and Results

Due to difficulty in obtaining the ground truth of the phone's orientation, the filter was tested with a combination of rotations around a fixed axis with constant speed where the ground truth was possible to estimate accurately, and longer rotations with shifting axes and speeds that ended at a known orientation. Each fixed axis rotation was repeated 3 times and the results averaged. The longer rotations were also performed 3 times.

Two other methods were used as a comparison. The first method was to calculate the orientation in global terms by normalizing the accelerometer readings to generate a 'down' unit vector and taking the component of the magnetometer that was not 'down' as north. This is the orientation method provided by the Android Application Programming Interface [1]. The second method was to integrate the gyroscope readings using the trapezoidal method. The orientation given with this method is with respect to the phone's initial position.

The state was initialized with calibrated gyroscope offsets, and the quaternion was initialized to [1 0 0 0]. The other state variables were initialized to 0.

	Approximate Error		
Test	**Accel. + Magnet. Only**	**Gyroscope Only**	**Kalman Filter**
360° rotation	8°	2°	2°
Random rotations (10 sec)	16°	10°	5°

The filter is able to improve on the two simpler methods for the longer test. Furthermore, unlike the 'gyroscope only' method, the orientation reported by the filter is with respect to the global coordinate system.

V. Conclusion and Future Work

The performance of orientation tracking can be improved over baseline methods using a Kalman Filter. A major limitation of the filter as implemented is its resource intensity. Running at 60hz, the filter takes up over 50% of the CPU on the Google Nexus S test platform, consuming more battery than desirable and leaving less CPU time for the application using the orientation tracking functionality. The resource intensity of the filter could be reduced by running it at a lower or variable update frequency, and also by computing the dynamics and measurement matrices using analytical rather than numerical derivatives.

One problem that may manifest itself when running the filter at a lower frequency is inaccuracy or divergence arising from the linearization process. If this does happen, one solution would be to implement an Unscented Kalman Filter, which has better performance for nonlinear systems.

VI. References

[1] Google Inc, "Android API Reference," Available: http://developer.android.com/reference/packages.html

[2] JMP van Waveren. "From Quaternion to Matrix and Back". Id Software, Inc. 2005.

[3] NOAA. "US/UK World Magnetic Model -- Epoch 2010.0" Available: http://www.ngdc.noaa.gov/geomag/WMM/DoDWMM.shtml

[4] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotic*s. Cambridge, Massachusetts: The MIT Press, 2005.