Linköping studies in science and technology. Thesis.
No. 1489

Topics in Localization and Mapping

Jonas Callmer

Division of Automatic Control
Department of Electrical Engineering
Linköping University, SE-581 83 Linköping, Sweden
http://www.control.isy.liu.se
callmer@isy.liu.se

Linköping 2011

This is a Swedish Licentiate's Thesis.

Swedish postgraduate education leads to a Doctor's degree and/or a Licentiate's degree.
A Doctor's Degree comprises 240 ECTS credits (4 years of full-time studies).
A Licentiate's degree comprises 120 ECTS credits,
of which at least 60 ECTS credits constitute a Licentiate's thesis.

Linköping studies in science and technology. Thesis.
No. 1489

Topics in Localization and Mapping

Jonas Callmer

callmer@isy.liu.se
www.control.isy.liu.se
Department of Electrical Engineering
Linköping University
SE-581 83 Linköping
Sweden

ISBN 978-91-7393-152-6 ISSN 0280-7971 LiU-TEK-LIC-2011:28

Copyright © 2011 Jonas Callmer

Printed by LiU-Tryck, Linköping, Sweden 2011

Till min familj

Abstract

The need to determine ones position is common and emerges in many different situations. Tracking soldiers or a robot moving in a building or aiding a tourist exploring a new city, all share the questions *"where is the unit?"* and *"where is the unit going?"*. This is known as the localization problem.

Particularly, the problem of determining ones position in a map while building the map at the same time, commonly known as the simultaneous localization and mapping problem (SLAM), has been widely studied. It has been performed in cities using different land bound vehicles, in rural environments using autonomous aerial vehicles and underwater for coral reef exploration. In this thesis it is studied how RADAR signals can be used to both position a naval surface vessel but also to simultaneously construct a map of the surrounding archipelago. The experimental data used was collected using a high speed naval patrol boat and covers roughly 32 km. A very accurate map was created using nothing but consecutive RADAR images.

A second contribution covers an entirely different problem but it has a solution that is very similar to the first one. Underwater sensors sensitive to magnetic field disturbances can be used to track ships. In this thesis, the sensor positions themselves are considered unknown and are estimated by tracking a friendly surface vessel with a known magnetic signature. Since each sensor can track the vessel, the sensor positions can be determined by relating them to the vessel trajectory. Simulations show that if the vessel is equipped with a global navigation satellite system, the sensor positions can be determined accurately.

There is a desire to localize firefighters while they are searching through a burning building. Knowing where they are would make their work more efficient and significantly safer. In this thesis a positioning system based on foot mounted inertial measurement units has been studied. When such a sensor is foot mounted, the available information increases dramatically since the foot stances can be detected and incorporated in the position estimate. The focus in this work has therefore been on the problem of stand still detection and a probabilistic framework for this has been developed. This system has been extensively investigated to determine its applicability during different movements and boot types. All in all, the stand still detection system works well but problems emerge when a very rigid boot is used or when the subject is crawling. The stand still detection framework was then included in a positioning framework that uses the detected stand stills to introduce zero velocity updates. The system was evaluated using localization experiments for which there was very accurate ground truth. It showed that the system provides good position estimates but that the estimated heading can be wrong, especially after quick sharp turns.

Populärvetenskaplig sammanfattning

Problemet att bestämma sin position uppstår i många olika sammanhang. Lokaliseringsproblemet, som det kallas, handlar i grunden om två fundamentala frågor; *"var är enheten?"* och *"vart är enheten på väg?"*. Lösningen är problemspecifik och baseras på sensorer och kunskaper om problemets natur.

Problemet dyker upp i många olika sammanhang. I flera fall, såsom positionsbestämning av en bil eller av ett fartyg i en skärgård, är ett globalt satellitbaserat positioneringssystem tillgängligt vilket gör problemet trivialt att lösa. I andra fall är det inte lika lätt. En robot som ska genomsöka ett avloppssystem, kartlägga ett korallrev eller utforska en okänd byggnad saknar möjligheten att bestämma sin position med hjälp av satelliter eftersom deras signaler är för svaga för att penetrera hus eller vatten. För dessa fall måste fristående lösningar utvecklas.

I flera sammanhang vill man inte bara positionera sin enhet utan samtidigt även skapa en karta över omgivningarna som man positionerar sig i. Detta kallas simultan lokalisering och kartskapning (SLAM) och är ett mycket populärt problem inom autonom robotik. SLAM har utförts i bland annat städer, i havet och i luften och anses idag vara ett väl utforskat problem. I denna avhandling studeras först problemet att bestämma ett fartygs position, kurs och hastighet samt även den omgivande skärgården genom att använda sekvenser av RADAR-signaler. Genom att matcha RADAR-bild mot RADAR-bild kan man estimera hur fartyget rör sig över tiden. Ett experimentdataset med RADAR-svep från en Stridsbåt 90 används för att utvärdera systemet. Experimentet är ungefär 32 km långt och visar att positioneringssystemet fungerar och att den karta som genereras motsvarar verkligheten mycket väl.

Ett andra bidrag handlar om ett problem som till synes ligger långt ifrån SLAM men vars lösning är mycket lik. Problemet handlar om att positionera ett antal sensorer utpositionerade på havsbottnen. Syftet med sensornätverket är att spana efter främmande fartyg genom att detektera de magnetfältsstörningar och ljud som dessa ger upphov till. Problemet är att sensorernas exakta positioner är okända eftersom strömmar och annat kan flytta på sensorerna medan de sjunker. Alternativet att förankra dem på en känd plats på botten är dyrt och tar tid. En lösning har därför tagit fram för att positionera sensorerna genom att låta dem spana på ett fartyg med känd magnetisk signatur. Exempelvis skulle det kunna vara samma fartyg som släppte ner sensorerna. Med hjälp av den magnetfältsstörning som fartyget ger upphov till när det åker runt i området så kan sensorernas positioner bestämmas. Simuleringar visar att systemet kan positionera sensorerna med hög noggrannhet, särskilt när fartygets rutt är loggad med ett satellitpositioneringssystem.

Inom vissa kretsar finns det en önskan om att kunna positionera personer som rör sig i en byggnad. Bland professionella användare såsom poliser eller brandkår tror man att ett system som visar var alla befinner sig skulle ge en stor förbättring av säkerheten för de inblandade. Det är exempelvis mycket enklare att hitta en skadad rökdykare om man vet hur byggnaden ser ut och var han är än om man

inte vet det. Tanken är att man ska sy in sensorer i deras uniformer för att med hjälp av dessa följa hur personerna rör sig i byggnaden. I vissa fall kan man anta att det finns kartor att tillgå (sjukhus, skolor, m.m.), i andra fall gör det inte det (lägenheter och villor). Positioneringssystemet måste fungera för alla typer av byggnader och scenarion och får inte baseras på orealistiska antaganden om exempelvis hur rökdykarna rör sig.

I denna avhandling studeras positioneringssystem för rökdykare som baseras på fotmonterade sensorer. Dessa sensorer mäter skons accelerationer och vinkelhastigheter vilket kan användas för att skatta dess position. Kan man dessutom detektera varje gång foten står stilla så kan positioneringsramverket uppdateras med informationen om att hastigheten är noll. Detta förbättrar positioneringsegenskaperna väldigt mycket. Ett probabilistiskt ramverk som detekterar att skon är stilla genom att studera accelerometerns och gyrots signaler har därför utvecklas. Systemet har genomgått omfattande tester för att utvärdera detektionsförmågan för olika rörelser, olika skor och olika sensorpositioner. En slutsats är att en väldigt kraftig känga gör att det är svårare att detektera när skon är still på marken när man genomför snabba rörelser såsom löpning. Anledningen är att skon är så kraftig att den inte omformar sig mot underlaget utan rullar mot det och därför aldrig är stilla. Det verkar också som att en fotmonterad sensor inte räcker till när man ska följa en krypande rökdykare. Systemet fungerar annars väl.

Stå still-detektionen har också använts i ett positioneringssystem som uppdateras med att hastigheten är noll varje gång foten nuddar backen. Systemet har utvärderats med hjälp av noggranna mätningar av skons sanna position. Experimenten visar att positionen kan skattas väl även om riktningen som rörelsen sker i ibland kan skattas fel, särskilt när stora, snabba svängar utförs. I framtiden kommer systemet förbättras genom att ta hänsyn till mer information såsom exempelvis kompassriktning.

Acknowledgments

First of all I would like to thank my supervisor Professor Fredrik Gustafsson for the excellent guidance and the inspiring projects you throw at me. Your insights, knowledge and efficiency never ceases to amaze me. My co-supervisor Dr David Törnqvist deserves a special thank for always being so enthusiastic, never being too busy and for putting that bit of pressure on me that I need to perform well. I can only hope that our collaborations continue to improve and become even more fruitful during the coming years. Also, thank you for proof reading the thesis.

Professor Lennart Ljung and Fredrik had the benevolence of inviting me to join the Automatic Control group. It has been a very interesting journey and I thank you for that. The group is skillfully headed by Professor Svante Gustafsson and Ninna Stensgård and before that by Lennart, Åsa Karmelind and Ulla Salaneck. You have managed to create an excellent group with supreme working conditions and I salute you for that.

My first contact with the world of academic research was the dynamic duo of Dr Juan Nieto and Dr Fabio Ramos of the ACFR in Sydney, Australia. You gave a fabulous introduction to the wonders of the academics with your fine wines and barbecues and let us not forget the *nomihodai* karaoke. I can only hope our paths in life will cross again in the near future.

I would also like to thank the rest of my colleagues in the Automatic Control group for our excellent working environment. I especially would like to point out the very collaborative atmosphere among the PhD students. We are all in this together and it makes this all go so much smoother. The makers if this LaTeX template, Dr Henrik Tidefelt and Dr Gustaf Hendeby, deserve a special thank for keeping the group's theses so beautiful.

Claes Norell and the rest of the Lambohov Fire Brigade have been kind enough to let me in on the secrets of firefighting. The input I get is invaluable since it allows me to always stay on the right track when developing their future positioning system. I thank you for this and hope that we can deepen our collaboration in the future.

Since life should not be all work and no play, *fackpampen* Lic Christian Lundquist, Lic Karl Granström, MSc Hanna Fager and MSc Martin Skoglund deserve a special mention. Our friendship, never ending discussions, *fikas*, travels and revolutions big and small make the times so much more fun. May it continue for a long time.

AFU (away from university), Krav Maga Nordic Linköping provides great physical workout and mental relaxation combined with a very cozy atmosphere. Think about anything else and you get a brand new face - there is no better motivation for focusing on what is right in front of you. Thank you Jens Berglund and Mårten Szymanowski for teaching me all those things I did not know could be done.

My old friends from the university, Lidingö, around the globe and *Broderskapet*

Jonas, I love that we take the time to bridge the geographical distances for never ending fun, travels and get-togethers. I hope we can keep it up well into the future since it means the world to me.

I would also like to thank my parents and my sister for your never ending love, support and encouragement despite my difficulties in explaining to you what I am actually doing. I would also like to take this opportunity to welcome the little one who will brighten our lives any day now. Being an uncle is going to be so much fun!

Last but not least, this work could never have been done without the financial support from the Strategic Research Center MOVIII, funded by the Swedish Foundation for Strategic Research, SSF and CADICS, a Linnaeus center funded by the Swedish Research Council. I would also like to acknowledge the research school Forum Securitatis in which I am participating.

Linköping, April 2011
Jonas Callmer

Contents

Notation **xv**

I Background

1 Introduction **3**
1.1 Problem Formulation 4
1.1.1 Indoor Localization 4
1.1.2 Surface Localization 4
1.1.3 Underwater Localization 5
1.2 Contributions 6
1.2.1 Additional Publications 6
1.3 Thesis Outline 7

2 Sensor Fusion **9**
2.1 Sensors 9
2.1.1 Inertial Measurement Unit 10
2.1.2 RADAR 11
2.1.3 Global Navigation Satellite System 12
2.1.4 VICON 13
2.2 Models 13
2.2.1 Continuous Models 13
2.2.2 Discrete Time Models 14
2.3 Estimation Theory 16
2.3.1 Kalman Filter 16
2.3.2 Extended Kalman Filter 17
2.3.3 Nonlinear Filtering Overview 18

3 Comparison between the Full and the Reduced Model **21**
3.1 Model Structures 21
3.1.1 Full Model Structure 22
3.1.2 Reduced Model Structure 23
3.2 Transfer Function Derivations 24

3.2.1 Transfer Function of Full Model 24
3.2.2 Transfer Function of Reduced Model 25
3.3 Bode Plots Evaluations 26
3.4 Discussion 28

4 Indoor Localization 31
4.1 Stand Still Detection 33
4.1.1 Related Work 33
4.1.2 Test Statistics Derivation 34
4.1.3 Test Statistic Distribution Validation 37
4.1.4 Hidden Markov Model 37
4.1.5 Experimental Results 41
4.1.6 Conclusions and Future Work 43
4.2 Stand Still Detection Performance for Different IMU Positions . . . 44
4.2.1 Experimental Results 45
4.2.2 Discussion 48
4.3 Localization Experiments 50
4.3.1 Measurements 50
4.3.2 Models 52
4.3.3 Filter 53
4.3.4 Experimental Results 54
4.3.5 Discussion and Future Work 55

5 Conclusions and Future Work 59
5.1 Conclusions 59
5.1.1 Indoor Localization 59
5.1.2 RADAR SLAM 60
5.1.3 Underwater Sensor Positioning 60
5.2 Future Work 60

A Quaternion Properties 63
A.1 Operations and Properties 63
A.2 Describing a Rotation using Quaternions 64
A.3 Rotation Matrix 64
A.4 Quaternion Dynamics 65

Bibliography 67

II Publications

A RADAR SLAM using Visual Features 73
1 Introduction 75
2 Background and Relation to SLAM 78
3 Theoretical Framework 79
3.1 Detection Model 79
3.2 Measurement Model 80

3.3 Motion Model 81
3.4 Multi-Rate Issues 82
3.5 Alternative Landmark Free Odometric Framework 83
4 SIFT Performance on RADAR Images 85
4.1 Matching for Movement Estimation 86
4.2 Loop Closure Matching 86
5 Experimental Results 87
5.1 Results 87
5.2 Map Estimate 90
6 Conclusions 90
Bibliography 92

B Silent Localization of Underwater Sensors using Magnetometers 95
1 Introduction 97
2 Methodology 99
2.1 System Description 99
2.2 State Estimation 101
2.3 Cramer-Rao Lower Bound 103
3 Simulation Results 104
3.1 Magnetometers Only 104
3.2 Magnetometers and GNSS 105
3.3 Trajectory Evaluation using CRLB 107
3.4 Sensitivity Analysis, Magnetic Dipole 107
3.5 Sensitivity Analysis, Sensor Orientation 108
4 Conclusions 110
Bibliography 112

Notation

Abbreviations

Abbreviation	Meaning
CRLB	Cramer-Rao Lower Bound
CRM	Corrosion Related Magnetism
EKF	Extended Kalman Filter
ESDF	Exactly Sparse Delayed-state Filter
FIM	Fisher Information Matrix
GNSS	Global Navigation Satellite Systems
GPS	Global Positioning System
HMM	Hidden Markov Model
IMU	Inertial Measurment Unit
PF	Particle Filter
RADAR	RAdio Detection And Ranging
RMSE	Root Mean Square Error
SIFT	Scale-Invariant Feature Transform
SLAM	Simultaneous Localization And Mapping
ZUPT	Zero Velocity Update

Estimation

Notation	Meaning
$\mathbf{x}_t$	State at time t
$\mathbf{y}_t$	Measurement at time t
T	Sampling Time
$\hat{\mathbf{x}}_{t+T\|t}$	State estimate at time $t+T$ given measurements up to and including time t
$P_{t+T\|t}$	Covariance of state estimate at time $t+T$ given measurements up to and including time t
p_t	Position state at time t
v_t	Velocity state at time t
a_t	Acceleration state at time t
q_t	Quaternion state at time t
ω_t	Angular velocity state at time t
w_t	Process noise at time t
r_t	Measurement noise at time t
Q	Process noise covariance
R	Measurement noise covariance

Part I

Background

1

Introduction

Localization is the problem of determining ones position. The problem ranges from robots to humans, from indoors to the oceans and from underground to the sky and further on into space. Estimating the position of a robot exploring a sewer system, determining the position of a ship in an archipelago or tracking a fire fighter searching through a burning building, all is localization.

If one is outdoors and a Global Navigation Satellite System (GNSS) like the Global Positioning System (GPS) is available, the positioning problem becomes trivial if the provided positioning accuracy is enough for the application. Measurements of ones position are then readily available making position estimation straightforward.

Unfortunately there are many environments where GNSS signals are not available. Indoors, underground or underwater the GNSS signals are too weak to be detected. Even outdoors the GNSS signals can be corrupted. This is commonly caused by reflections against houses or that the lines of sight to the satellites are blocked by trees or high rise buildings. Lately, intentional or unintentional jamming of the GPS signals has also emerged as a potential major problem. All this makes systems depending entirely on GNSS vulnerable.

Different, redundant means of positioning are therefore required, ones that are tailored for each specific problem. The solutions must be reliable and use all other available information to get the best possible positioning estimate. That is the problem of localization.

1.1 Problem Formulation

Three problems have been studied in this thesis. Even though they may seem different at first, they all share the problem of localization. The first is indoor localization for firefighters and the other two cover maritime localization problems.

1.1.1 Indoor Localization

The problem of indoor localization for professional users has received a lot of attention lately since it is a fundamental problem in a multitude of situations; Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007); Godha et al. (2006); Woodman and Harle (2009); Grzonka et al. (2010); Aggarwal et al. (2011); Jiménez et al. (2010); Robertson et al. (2009); Widyawan et al. (2008). Being able to track the position of each individual firefighter, police officer or soldier moving around in a building is the dream of every operational management. In case something urgent happens, knowing exactly where all personnel are and where they have been enables swift and accurate cooperation to solve the problem. Having a positioning system would therefore greatly enhance the safety of the personnel. The problem of positioning firefighters in a smoke filled house and the envisioned position presentation is shown in Figure 1.1.

There is a tendency to view all three users; firefighters, police officers and soldiers, as having the same needs for the same problem and therefore require the same solution. We do not fully see it that way. The level of stress is different and the tactics of how to search a building are entirely different. This makes the working conditions different enabling some solutions to work in some situations but not in others. In order to achieve acceptable positioning all must be taken into consideration and exploited - including tactics.

It is the purpose of this thesis and our future work within the area to find a solution to the indoor localization problem for firefighters entering a building of limited size, like a residential house or a small office, of which there is no previous knowledge about the layout of the building. The solution should be as simple as possible, using as few sensors as possible and based on as few assumptions about user movements and the environment as possible.

1.1.2 Surface Localization

Modern maritime navigation is very GPS centered. The system is not only used for positioning but often also as a compass, to track communications satellites or as a very accurate time measurement. Since the GPS signals are easily jammed, a backup system is needed when navigating in critical environments. We present a solution based entirely on the measurements from the ship's RADAR where the scans are used to estimate the position, velocity and heading of the vessel.

(a) Firefighters getting prepared to enter a burning building.

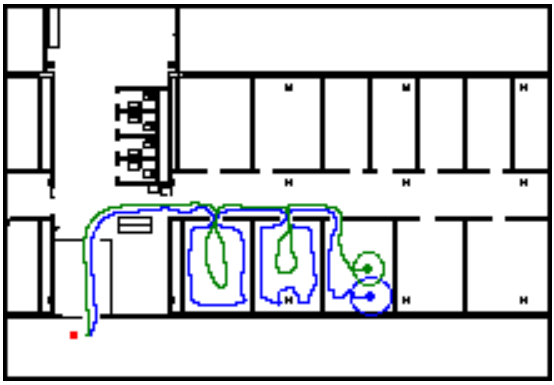

(b) Presentation of the current and previous positions of the firefighters with uncertainties.

Figure 1.1: *When firefighters enter a burning building, left, they are equipped with sensors. Their estimated positions are presented to the operation management, right.*

1.1.3 Underwater Localization

A passive surveillance sensor network can track a passing vessel using underwater sensors that sense the magnetic field disturbances and noises caused by the vessel. The problem is that the exact sensor positions are seldom known unless a large amount of time and money has been spent on determining their exact positions. Currents also cause the sensors to move while sinking making rapid sensor deployment difficult. Without correct sensor positions, accurate tracking of enemy vessels becomes impossible. In this thesis we have studied the problem of determining the positions of the sensors using their sensor readings when a friendly vessel with a known magnetic signature is passing by. By knowing where the vessel has been, the positions of the sensors can be accurately determined. Now when the true sensor positions are known the network can start performing its original task; searching for naval intruders.

1.2 Contributions

The main contributions in this thesis are

- A probabilistic framework for stand still detection for a foot mounted IMU.
- A RADAR based backup system for positioning a surface vessel in case of GPS outage.
- A sensor positioning framework where the movements of a friendly vessel are used by the sensor network to determine their individual positions.

Published work of relevance to this thesis are listed below.

> J. Callmer, D. Törnqvist, H. Svensson, P. Carlbom, and F. Gustafsson. Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2010c. Under revision.
>
> J. Callmer, M. Skoglund, and F. Gustafsson. Silent localization of underwater sensors using magnetometers. *EURASIP Journal on Advances in Signal Processing*, 2010a.
>
> J. Callmer, D. Törnqvist, and F. Gustafsson. Probabilistic stand still detection using foot mounted IMU. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010b.
>
> J. Rantakokko, P. Strömbäck, J. Rydell, J. Callmer, D. Törnqvist, F. Gustafsson, P. Händel, M. Jobs, and M. Grudén. Accurate and reliable soldier and first responder indoor positioning: Multi-sensor systems and cooperative localization. *IEEE Wireless Communications Magazine*, 2011. Accepted for publication.

1.2.1 Additional Publications

Other publications where the author has contributed that are not covered in this thesis are listed and briefly described below.

Two publications are about using a three axis magnetometer to track vehicles using the magnetic disturbances induced by the vehicles. The publications are based on a master thesis project undertaken by Niklas Wahlström that was supervised by the author.

> N. Wahlström, J. Callmer, and F. Gustafsson. Single target tracking using vector magnetometers. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011.
>
> N. Wahlström, J. Callmer, and F. Gustafsson. Magnetometers for tracking metallic targets. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010.

As backup for GNSS, Lindsten et al. (2010) covered the problem of using preexisting maps and environmental classification to create a measurement of the global

position of an Unmanned Aerial Vehicle (UAV). Photos from a downwards facing camera on the UAV were classified into grass, houses, roads etc which were matched to a map of the area. The main contribution of the author was within the classification parts.

> F. Lindsten, J. Callmer, H. Ohlsson, D. Törnqvist, T. B. Schön, and F. Gustafsson. Geo-referencing for UAV navigation using environmental classification. In *Proceedings of 2010 International Conference on Robotics and Automation (ICRA)*, 2010.

The last two publications are spinoffs from the master thesis project undertaken by Karl Granström and the author in 2007/08. The papers are about loop closure detection in urban SLAM by matching laser scans in Granström et al. (2009) and photos in Callmer et al. (2008). The aim was to find reliable loop closure detection methods that worked in large scale problems.

> K. Granström, J. Callmer, F. Ramos, and J. Nieto. Learning to detect loop closure from range data. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2009.

> J. Callmer, K. Granström, J. Nieto, and F. Ramos. Tree of words for visual loop closure detection in urban slam. In *Proceedings of the 2008 Australasian Conference on Robotics and Automation (ACRA)*, 2008.

1.3 Thesis Outline

The thesis is divided into three parts. The first one starting on Chapter 2 is a brief introduction to the sensor fusion tools of modelling, sensors and estimation theory that are necessary in a localization framework. Chapter 3 dives deeper into the modelling problem with a discussion about the differences between two modelling approaches. The second part is about indoor navigation for professional users and is covered in Chapter 4. A framework for stand still detection using a foot mounted inertial measurement unit is presented and thoroughly evaluated. The framework is thereafter included in localization experiments with zero velocity updates. The experiments were performed in the VICON lab providing extremely accurate ground truth data which is very unusual to have for these types of experiments, if not unique. Chapter 5 summarizes this part of the thesis with conclusions and a discussion about future work.

The third part consists of the publications; Paper A covers Simultaneous Localization and Mapping (SLAM) in an archipelago using the RADAR sensor of a high speed naval patrol boat. The second publication, Paper B, is about determining the positions of a large number of underwater sensors equipped with magnetometers and microphones. By tracking the movements of a friendly vessel with a known magnetic signature, the sensor positions can be determined.

2

Sensor Fusion

Sensor fusion is the problem of estimating some properties x_t of a unit using sensors that provide measurements y_t that depend on those properties. In order to do this, models of how y_t are related to x_t and models of how x_t change over time, are used.

The properties x_t are called states and can represent any sought system property. The states can for example be related to the sensor platform representing the position or orientation of the unit, they can be some unknown properties of the unit that is needed but unknown such as its weight or they can be related to the surrounding environment such as the positions of some environmental landmarks, among others. In the problem of localization the sought after states are commonly position, velocity and orientation answering the fundamental questions ”where is the unit?” and ”where is the unit going?”.

The states are estimated using a filter that fuses the information from the sensors and the models. Besides from the estimated states the filter also provides the uncertainties of those state estimates. A schematic overview of a sensor fusion framework is given in Figure 2.1.

2.1 Sensors

The sensors used in localization are usually of two kinds. Either they measure the movements of the unit or they measure some aspect of the surroundings. The former could for example be accelerometers measuring the unit accelerations, gyros measuring its angular velocities or if it is a robot, wheel encoders that measure how much the wheels are rotating. The latter could be cameras filming the surroundings, RADARs or laser range finders measuring the distances to objects

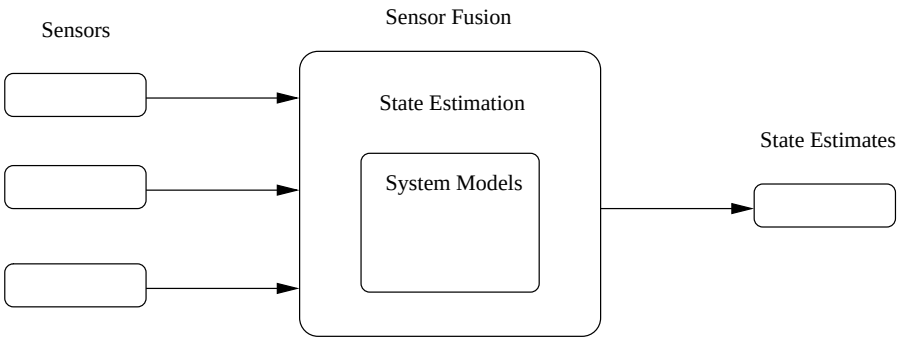

Figure 2.1: *Overview of a sensor fusion framework*

around the unit, or magnetometers measuring the surrounding magnetic field, among others.

The sensors used in this thesis are an inertial measurement unit in the indoor localization experiments and a RADAR sensor in the maritime RADAR localization experiments. The underwater sensor positioning was only performed using simulated magnetometer and microphone data why no sensors were used.

2.1.1 Inertial Measurement Unit

An inertial measurement unit (IMU) contains an accelerometer and a gyroscope measuring the accelerations and angular velocities of the unit. These are measured in three dimensions making it theoretically possible to track the orientation of the user by integrating the angular velocity measurements. If the orientation is known, the direction of down is known making it possible to remove the gravity component from the acceleration measurements. Double integrating the remaining user accelerations gives the position of the sensor, in theory that is. In practice the noisiness of the signals makes it only possible to accurately track the sensor position and orientation over a short time interval.

The sensor unit is commonly also equipped with a magnetometer, measuring the magnetic field around the sensor. When outdoors, the earth magnetic field is commonly free from magnetic inference. Indoors, steel structures, electrical wiring and cabinets produce severe magnetic disturbances making the magnetometer more unreliable.

When the IMU is at rest, a reading of the uncorrupted earth magnetic field and the accelerometer can be used to determine the orientation of the sensor unit. The accelerometer gives the direction of up and the magnetometer effectively is a compass giving the direction of north. The orientation can thereafter be tracked during movements over a short time interval using the gyros.

The IMU used in the indoor localization related experiments in Chapter 4 is an Xsens MT motion sensor, Figure 2.2, equipped with a three axis accelerometer, a three axis gyro and a three axis magnetometer. The sensor is also equipped with

Figure 2.2: *An Xsens MT motion sensor, courtesy of Xsens Technologies B.V.*

a thermometer to compensate for the temperature dependency of the different components. The signals were sampled in 100 Hz using a 16 bit A/D converter.

The gyroscope and accelerometer are based on *micro-machined electromechanical systems* (MEMS) technology making it small, light, cheap and low on power consumption. The main drawback is their quite poor performance even if the sensors have been thoroughly calibrated. Gyro biases, scaling errors and temperature dependent disturbances often remain.

2.1.2 RADAR

A pulse RADAR (RAdio Detection And Ranging) sends out radio waves in different directions which are reflected or scattered when hitting an object. The reflected signals are picked up by a receiver, usually at the same place as the transmitter, and the time of flight for the signal is calculated. This time is proportional to the distance to the object that reflected the signal and the heading of the sensor when the signal was transmitted gives the direction to the object. The strength of the reflected signal can also provide some information about the properties of the object.

Two measurements are provided by each reflecting object; range and angle from the sensor to the object.

$$r_t = \sqrt{(s_{x,t} - p_{x,t})^2 + (s_{y,t} - p_{y,t})^2} \tag{2.1}$$

$$\alpha_t = \arctan \frac{s_{y,t} - p_{y,t}}{s_{x,t} - p_{x,t}} \tag{2.2}$$

where $s = (s_{x,t}, s_{y,t})$ is the position of the reflecting structure and $p = (p_{x,t}, p_{y,t})$ is the position of the RADAR sensor at time instant t. Since the uncertainties in angle and range are independent, the total measurement uncertainties will be banana shaped.

A naval RADAR commonly rotates with a constant speed, transmitting and receiving in one direction at a time. The reflections are plotted in the current direction

when they are received. This gives a circular image of the surrounding islands and vessels that is updated one degree at a time. By saving one 360° RADAR sweep as an image, a view of the surroundings is provided.

The RADAR sensor used in Paper A was a military one making the characteristics of that sensor secret. What we do know is that it had a range of roughly 5 km and a range resolution of about 5 meters. It rotates one revolution in 1.5 seconds giving measurements in roughly 2000 directions.

One way of using a RADAR in localization is to take some strong reflections in a full 360° RADAR scan and try to detect them again in the next scan. The objects creating these reflections are called landmarks and are assumed stationary. By measuring the distance and heading to the landmarks and see how these change over time, an estimate of how the RADAR equipped unit is moving can be made. If some landmarks move in a manner that is inconsistent with the other landmarks, it is probably a different unit and the reflections should not be used for localization.

2.1.3 Global Navigation Satellite System

GNSS systems use multiple satellites and triangulation to determine the position of the user. The most well known system, GPS, provides a positioning accuracy of about 10 meters. Besides from location, it also gives accurate time information making it useful also in applications where only accurate time and not position is needed, for example in cellphone base station synchronization. The system consists of 30 satellites and free line of sight to at least 4 of them is required for the positioning to work.

Other systems do exist or are planned. The Russian GLONASS system mostly covers the northern hemisphere, in particular Russia, and is today short of the 24 satellites needed to cover the whole planet. The European Galileo system will use 30 satellites to cover the entire planet and is scheduled to be operational around 2014. Also a Chinese system, COMPASS, using 35 satellites to cover the planet will be deployed in the future. As of today, a smaller system covering only China and the immediate surroundings is in place. A future GNSS receiver, using signals from all systems will pretty much always have free line of sight to at least 4 satellites. This will give accurate positioning also in places that are difficult to cover today such as urban canyons.

One shortcoming with GNSS systems is the weakness of the signals. The signal is weaker than the background noise and only because the receivers know what to look for can the signals be found. This makes the system sensitive to signal disturbances due to intentional or unintentional jamming. Today, GPS jammers that can easily knock out all GPS reception in an area of many square kilometers are available at a low cost; Economist (2011); Grant et al. (2009). This problem and a suggested solution for maritime vessels is discussed in more detail in Paper A.

2.1.4 VICON

The VICON[1] lab at LiU has been used to collect ground truth for the positioning experiments. The VICON system uses 10 infrared cameras and infrared lamps to track reflective balls that are attached to the object of interest. In the indoor positioning experiments, the balls were attached to the boot making the true movements of the IMU trackable in real time. The positioning precision provided by the system is on millimeter level. The size of the active area in the lab is about 3 by 5 meters forcing the experiments to traverse the same area multiple times in order to make them large enough.

The author would like to thank Piotr Rudol and the rest of the UAS Technologies Lab, Artificial Intelligence and Integrated Computer Systems Division (AIICS)[2] at the Department of Computer and Information Science (IDA) at LiU for their assistance during these experiments.

2.2 Models

To compute how a system is behaving, a computer needs mathematical models that describe the properties of the system. Mathematical models are most commonly described using differential equations but to make the models more useful for computers, approximate time discreet difference models are needed.

2.2.1 Continuous Models

The models are commonly on a state space form where a state vector x_t describes the system properties at time t. The states are related to the dynamical properties of the system and also to measurements y_t. The fundamental time continuous model is

$$\dot{x}_t = f(x_t, u_t, w_t) \tag{2.3}$$

$$y_t = h(x_t, u_t, v_t) \tag{2.4}$$

where u_t is a known input signal and w and v are noise terms illustrating the model errors. $\mathrm{f}(\cdot)$ and $\mathrm{h}(\cdot)$ are in general nonlinear functions relating the dynamic properties and the measurements to the system states.

These two types of models are fundamental in a sensor fusion framework. Dynamical models (2.3) describe how the unit states change over time. These are both used to restrict the type of possible movements of the unit and to physically relate the states to eachother. The former can for example be a vehicle model that states that a vehicle moves forward or backward, not sideways. The latter is that velocity estimates over time translate into changes in estimated position, among others. Since each physical model is a simplification of the true system the dynamical model is also associated with a process noise w that reflects the un-

[1] http://www.vicon.com
[2] http://www.ida.liu.se/divisions/aiics/aiicssite/index.en.shtml

certainties in the model and also the fundamental simplifications of the system which have been imposed.

Measurement models (2.4) relate the sensor measurements to the unit states. This can for example be range and angle measurements that are related to position states in the state vector. Related to each measurement is a measurement noise ν representing the ever present measurement noise.

2.2.2 Discrete Time Models

Since continuous models cannot be used by a computer they must be discretized. For most cases discretization is a complex task. One exception is a linear system

$$\begin{aligned} \dot{x}_t &= Ax_t + Bw_t \\ y_t &= Cx_t + \nu_t \end{aligned} \tag{2.5}$$

where the process noise w_t is assumed piecewise constant over the sampling interval T. The matrices of the sampled systems can then be computed as

$$F = e^{AT} \tag{2.6}$$

$$G = \int_0^T e^{A\tau} d\tau B \tag{2.7}$$

giving the time discrete linear system

$$\begin{aligned} x_{t+T} &= Fx_t + Gw_t \\ y_t &= Cx_t + D\nu_t. \end{aligned} \tag{2.8}$$

For most other cases sampling a continuous system is quite challenging. For details see Gustafsson (2010).

A general description of a physical system as a state space model in discrete time is

$$\begin{aligned} x_{t+1} &= f(x_t, w_t) \\ y_t &= h(x_t, \nu_t). \end{aligned} \tag{2.9}$$

An important special case is when the process and measurement noises are modeled as additive

$$\begin{aligned} x_{t+1} &= f(x_t) + w_t \\ y_t &= h(x_t) + \nu_t. \end{aligned} \tag{2.10}$$

It is an intuitively straightforward model with a deterministic part utilizing basic physical properties and a random part representing everything that cannot be prophesied that still affects the system. An example of a system with partially nonlinear dynamics and measurements is given in Example 2.1.

The linear system (2.8) is an important special case of modelling since it allows Kalman filter theory to be applied when solving the problem if the noises are assumed Gaussian.

2.1 Example

Model of an inertial navigation system estimating the sensor position using an accelerometer and a gyro and also global measurements of position from for example a GPS receiver.

The states are position p, velocity v and acceleration a, all in global coordinates, followed by orientation in quaternions q and angular velocity ω in the local coordinate system. The measurements are acceleration y_a and angular velocity y_ω, both measured in the local coordinate system and global position y_p from the GPS receiver. All measurements have additive noise r. The model assumes constant acceleration and angular velocity. Since this is not true the model is slightly wrong. Process noise w has therefore been added to represent the model errors that are introduced by these assumptions.

Quaternion dynamics and properties are described in Appendix A but a very brief explanation will be given here. $S'(q_t)\omega_t$ describes how local angular velocities translate into changes in quaternions. $R(q_t)$ is the rotation matrix from global to local coordinate systems based on the quaternions q_t.

$$\begin{pmatrix} p_{t+1} \\ v_{t+1} \\ a_{t+1} \\ q_{t+1} \\ \omega_{t+1} \end{pmatrix} = \begin{pmatrix} I & TI & \frac{T^2}{2}I & 0 & 0 \\ 0 & I & TI & 0 & 0 \\ 0 & 0 & I & 0 & 0 \\ 0 & 0 & 0 & I & \frac{T}{2}S'(q_t) \\ 0 & 0 & 0 & 0 & I \end{pmatrix} \begin{pmatrix} p_t \\ v_t \\ a_t \\ q_t \\ \omega_t \end{pmatrix} + \begin{pmatrix} \frac{T^3}{6}I & 0 \\ \frac{T^2}{2}I & 0 \\ TI & 0 \\ 0 & \frac{T^2}{2}S'(q_t) \\ 0 & T \end{pmatrix} \begin{pmatrix} w_a \\ w_\omega \end{pmatrix} \tag{2.11}$$

$$\begin{pmatrix} y_{a,t} \\ y_{\omega,t} \\ y_{p,t} \end{pmatrix} = \begin{pmatrix} 0 & 0 & R(q_t) & 0 & 0 \\ 0 & 0 & 0 & 0 & I \\ I & 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} p_t \\ v_t \\ a_t \\ q_t \\ \omega_t \end{pmatrix} + \begin{pmatrix} R(q_t)g \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} r_a \\ r_\omega \\ r_p \end{pmatrix} \tag{2.12}$$

The acceleration measurements y_a contain both user accelerations and the (inverse) gravity component g why it must be included in the measurement model.

As an alternative, a reduced model form is often used. The acceleration and gyro measurements are used as input signals in the dynamical model removing the need of states a_t and ω_t. That leaves us with only the position measurement in the measurement model.

$$\begin{pmatrix} p_{t+1} \\ v_{t+1} \\ q_{t+1} \end{pmatrix} = \begin{pmatrix} I & TI & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix} \begin{pmatrix} p_t \\ v_t \\ q_t \end{pmatrix} + \begin{pmatrix} \frac{T^2}{2}I & 0 \\ TI & 0 \\ 0 & \frac{T}{2}S'(q_t) \end{pmatrix} \begin{pmatrix} R^T(q_t)y_{a,t} - g + R^T(q_t)r_a + w_a \\ y_{\omega,t} + r_\omega + w_\omega \end{pmatrix} \tag{2.13}$$

$$y_p = p_t + r_p \tag{2.14}$$

What are the practical differences between these two approaches? This will be studied in Chapter 3.

2.3 Estimation Theory

The estimation problem is the problem of estimating the posterior distribution of the states given the measurements, $p(x_t|y_{1:t})$. The states are often intricately related to the measurements, making them difficult to estimate. With the use of Bayes' theorem

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)} \tag{2.15}$$

the problem can be reformulated into three straightforward parts. $p(y|x)$ is the likelihood of receiving the measurement y given the states x are true, $p(x)$ is the prior probability of the states incorporating all our previous knowledge about the states and $p(y)$, the probability of the measurement, normalizes the state probabilities.

Estimation theory can be divided into two cases; linear and nonlinear estimation. Linear estimation is straightforward and the results are trustworthy but the problem is in reality rare. Nonlinear estimation is more difficult and the solutions are prone to diverge but unfortunately the problem is very common.

In this section we will present the Kalman filter used for linear estimation problems, the extended Kalman filter used for slightly nonlinear problems and then briefly discuss some additional nonlinear filter methods that are not used in this thesis.

2.3.1 Kalman Filter

The linear estimation problem where a linear system (2.8) is assumed to have Gaussian noise is optimally solved using the Kalman filter, Kalman (1960). Since the problem is Gaussian, estimating the mean and the covariance of the states provides the entire solution to $p(x_t|y_{1:t})$.

The Kalman filter works in a two step procedure with a time update and a measurement update. The time update predicts the future states $\hat{x}_{t+1|t}$ using the dynamic model. Since the model is not perfect, the process noise covariance Q is added to the state covariance P to illustrate the increase in uncertainty.

$$Q = \mathrm{Cov}(w_t) \tag{2.16}$$

The measurement update uses the difference between the measurement and the predicted measurement, the innovation, to update the states. How much the new measurement should affect the states is decided by the Kalman gain, K. K is depending on Q and a measurement noise term R which describes how trustworthy the measurements are. The relation between these two parameters determines the filter performance. If Q is small in relation to R, the model is deemed more reliable than the measurements and vice versa. The ratio between Q and R affects the state estimates while the magnitudes of Q and R determines the size of the state covariances.

The equations defining the Kalman filter are shown in Algorithm 1.

Algorithm 1 Kalman Filter

Require: Signal model (2.8), initial state estimate $\hat{x}_{0|0}$ and covariance $P_{0|0}$.

1: Time Update

$$\begin{aligned} \hat{x}_{t+1|t} &= A\hat{x}_{t|t} \\ P_{t+1|t} &= AP_{t|t}A^T + Q \end{aligned} \tag{2.17}$$

2: Measurement Update

$$\begin{aligned} K_{t+1} &= P_{t+1|t}C^T\left(CP_{t+1|t}C^T + R\right)^{-1} \\ \hat{x}_{t+1|t+1} &= \hat{x}_{t+1|t} + K_{t+1}\left(y_{t+1} - C\hat{x}_{t+1|t}\right) \\ P_{t+1|t+1} &= P_{t+1|t} - K_{t+1}CP_{t+1|t} \end{aligned} \tag{2.18}$$

2.3.2 Extended Kalman Filter

If the problem is of the form (2.9) or (2.10) and only mildly nonlinear, the extended Kalman filter (EKF) can be applied. It approximates the nonlinearities using a first order Taylor approximation around the latest state estimate and then applies the Kalman filter to the linearized problem. Convergence cannot be guaranteed, particularly since the EKF gives the optimal solution to the wrong problem. That means, the solution to the linearized problem is the optimal one, unfortunately the linearized problem is not the true problem. Despite this, the EKF most often works well.

Starting with the nonlinear function (2.10), a first order Taylor expansion of the measurement function $\mathrm{h}(\cdot)$ around the linearization point $\hat{x}$ is

$$h(x_t) \approx h(\hat{x}_t) + h'_x(\hat{x}_t)(x_t - \hat{x}_t) \tag{2.19}$$

where $h'_x(\hat{x}_t)$ is the Jacobian

$$h'_x(\hat{x}_t) = \left.\frac{\partial h(x_t)}{\partial x_t}\right|_{x_t=\hat{x}_t}. \tag{2.20}$$

The measurement model can now be reformulated according to

$$\begin{aligned} y_t &= h(x_t) + v_t \Leftrightarrow \\ y_t - h(\hat{x}_t) + h'_x(\hat{x}_t)\hat{x}_t &= h'_x(\hat{x}_t)x_t + v_t \Leftrightarrow \\ \bar{y}_t &= h'_x(\hat{x}_t)x_t + v_t. \end{aligned} \tag{2.21}$$

Correspondingly the dynamical model can be expanded around $\hat{x}_t$ giving the new model

$$\begin{aligned} x_{t+1} &= f(x_t) + w_t \Leftrightarrow \\ x_{t+1} - f(\hat{x}_t) + f'_x(\hat{x}_t)\hat{x}_t &= f'_x(\hat{x}_t)x_t + w_t \Leftrightarrow \\ \bar{x}_{t+1} &= f'_x(\hat{x}_t)x_t + w_t \end{aligned} \tag{2.22}$$

where

$$f_x'(\hat{x}_t) = \left.\frac{\partial f(x_t)}{\partial x_t}\right|_{x_t=\hat{x}_t}. \tag{2.23}$$

If the signal model is (2.9), the Jacobians with respect to the noise parameters w and v are also needed.

$$f_w'(\hat{x}_t) = \left.\frac{\partial f(x_t, w_t)}{\partial w_t}\right|_{x_t=\hat{x}_t} \qquad h_v'(\hat{x}_t) = \left.\frac{\partial h(x_t, v_t)}{\partial v_t}\right|_{x_t=\hat{x}_t} \tag{2.24}$$

The extended Kalman filter is summarized in Algorithm 2.

Algorithm 2 Extended Kalman Filter

Require: Signal model (2.9), initial state estimate $\hat{x}_{0|0}$ and covariance $P_{0|0}$.

1: Time Update

$$\begin{aligned} \hat{x}_{t+1|t} &= f(\hat{x}_{t|t}) \\ P_{t+1|t} &= f_x'(\hat{x}_{t|t})P_{t|t}f_x'(\hat{x}_{t|t})^T + f_w'(\hat{x}_{t|t})Qf_w'(\hat{x}_{t|t})^T \end{aligned} \tag{2.25}$$

2: Measurement Update

$$\begin{aligned} S_{t+1} &= h_x'(\hat{x}_{t+1|t})P_{t+1|t}h_x'(\hat{x}_{t+1|t})^T + h_v'(\hat{x}_{t+1|t})Rh_v'(\hat{x}_{t+1|t})^T \\ K_{t+1} &= P_{t+1|t}h_x'(\hat{x}_{t+1|t})^T S_{t+1}^{-1} \\ \hat{x}_{t+1|t+1} &= \hat{x}_{t+1|t} + K_{t+1}\Big(y_{t+1} - h(\hat{x}_{t+1|t})\Big) \\ P_{t+1|t+1} &= P_{t+1|t} - P_{t+1|t}h_x'(\hat{x}_{t+1|t})^T S_{t+1}^{-1} h_x'(\hat{x}_{t+1|t})P_{t+1|t} \end{aligned} \tag{2.26}$$

2.3.3 Nonlinear Filtering Overview

For very nonlinear filtering problems, the EKF cannot be used. Other solutions exist, each with its pros and cons. Below follows a brief description of some of the most common nonlinear filters.

The particle filter does not try to model the entire pdf of the posterior distribution of x_t, Gordon et al. (1993). Instead it approximates the distribution using N samples, so called particles, each having a weight $\{\gamma_t^i\}_{i=1}^N$. N has to be large to make the approximation valid. The resulting approximation of the posterior distribution is $p(x_t|y_{1:t}) \approx \sum_{i=1}^N \gamma_t^i \delta(x_t - x_t^i)$, $\delta(\cdot)$ being Dirac's delta function. The particle filter can be used for very nonlinear estimation problems but it is not guaranteed to converge. Still, the filter has become very popular since it can handle difficult nonlinear non-Gaussian problems, is easy to implement and is quite intuitive.

If the dimension of the problem is low, the Point Mass Filter can be used. It grids the state space and computes the posterior distribution in all the grid points. It can be applied to any nonlinear and non-Gaussian model. Its drawback is that the number of grid points grows exponentially with the number of dimensions of the state vector making the filter only applicable for low dimensional problems.

Further, the unscented Kalman filter, Julier et al. (1995), deterministically samples so called sigma points around the mean of the state distribution. These sigma points are propagated through the nonlinear models and are used to reconstruct the mean and covariance of the state estimates. This will more accurately describe the true distribution of the states compared to EKF and requires no time consuming Jacobian computations.

Other nonlinear filtering methods such as Gaussian sum Kalman filters also exist. It models the state pdf using multiple Gaussians providing multimodal applicability.

3

Comparison between the Full and the Reduced Model

A common situation in navigation systems is that we want to estimate a state but we only have the derivative of the state available as a measurement. One example is when heading is estimated using a gyro, another is when velocity or position is estimated using an accelerometer. The situation can also be as in Example 2.1 in Section 2.2.2 where both position and acceleration measurements are available.

When we choose our dynamical model, the question is if we should model the measurements as outputs or reduce the state space by treating the measurements as inputs. In the former, all measurements have a corresponding state, in the case of an acceleration measurement we have an acceleration state as well as a velocity state and a position state. This was the form described initially in Example 2.1. In the latter, the derivative measurements are used as an input signal during the time update of the integral states. In this case the acceleration measurements are only used as time update inputs for the velocity and position states. This was the second model form described in Example 2.1. This reduced model form is also known as a Luenberger observer, Rugh (1996).

Since the reduced input model has fewer states than the full state model it requires fewer computations but what are the other differences between these model structures? That will be studied in this section.

3.1 Model Structures

One way to study the differences between the model structures is to compare the transfer functions of the systems. Unfortunately, it is not possible to calculate the transfer function from an acceleration measurement to a position estimate if no position measurement is available since the estimation uncertainty will grow in-

definitely with time. This is natural since the problem is not observable. If both position and acceleration measurements are available, the position can be estimated using a stationary Kalman filter. Since the filter is stationary, the transfer function can be calculated. By choosing the measurement noise of the position measurement as extremely large, we are in practice back to the structure of estimating a position using only acceleration measurements.

For the analysis, the following model states and measurements are used. x_p, x_v and x_a represent the position, velocity and acceleration estimates, respectively, and y_p, y_v and y_a represent the position, velocity and acceleration measurements, respectively. Two transfer functions are interesting to investigate; acceleration measurement to velocity estimate and acceleration measurement to position estimate. Note that the former could equally well be seen as any integration, for example from angular velocity measurements to heading estimates.

3.1.1 Full Model Structure

In the full model all states including the acceleration are estimated. This gives the state and measurement vectors

$$\mathbf{x}_t = \begin{pmatrix} x_{p,t} & x_{v,t} & x_{a,t} \end{pmatrix}^T$$
$$\mathbf{y}_t = \begin{pmatrix} y_{p,t} & y_{v,t} & y_{a,t} \end{pmatrix}^T.$$

The basic full model structure can be written as

$$\begin{aligned} \mathbf{x}_{t+T} &= A\mathbf{x}_t + Bw_t \\ \mathbf{y}_t &= C\mathbf{x}_t + \mathbf{r}_t \end{aligned} \tag{3.1}$$

where $\mathbf{r}_t$ and w_t are measurement noises and process noise, respectively.

The full model is in this case

$$\begin{aligned} \begin{pmatrix} x_{p,t+T} \\ x_{v,t+T} \\ x_{a,t+T} \end{pmatrix} &= \begin{pmatrix} 1 & T & \frac{T^2}{2} \\ 0 & 1 & T \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_{p,t} \\ x_{v,t} \\ x_{a,t} \end{pmatrix} + \begin{pmatrix} \frac{T^3}{6} \\ \frac{T^2}{2} \\ T \end{pmatrix} w_t \\ \begin{pmatrix} y_{p,t} \\ y_{v,t} \\ y_{a,t} \end{pmatrix} &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_{p,t} \\ x_{v,t} \\ x_{a,t} \end{pmatrix} + \begin{pmatrix} r_{p,t} \\ r_{v,t} \\ r_{a,t} \end{pmatrix} \end{aligned} \tag{3.2}$$

where the process noise $w \sim N(0, \sigma_w^2)$ and $r_p \sim N(0, \sigma_p^2)$, $r_v \sim N(0, \sigma_v^2)$ and $r_a \sim N(0, \sigma_a^2)$ are noises from the position, velocity and acceleration measurements, respectively. By choosing σ_p^2 and σ_v^2 very large, they are deemed insignificant by the filter and only the acceleration measurements are used to estimate the velocity and position.

One fundamental property of the full model structure is that it assumes that the acceleration is constant over the sampling interval.

3.1.2 Reduced Model Structure

In the reduced model the acceleration measurements are used as direct inputs in the time updates of the velocity and position estimates. This model also assumes that the acceleration is constant over the sampling interval.

Again, measurements of position and velocity are needed to get a stationary Kalman filter. The state and measurement vectors used are

$$\begin{aligned}\mathbf{x}_t^u &= \begin{pmatrix} x_{p,t} & x_{v,t} \end{pmatrix}^T \\ \mathbf{y}_t^u &= \begin{pmatrix} y_{p,t} & y_{v,t} \end{pmatrix}^T .\end{aligned} \tag{3.3}$$

The superscript u will be used throughout the chapter to indicate that we mean the reduced model.

The basic model structure is in this case

$$\begin{aligned}\mathbf{x}_{t+T}^u &= A^u \mathbf{x}_t^u + B_1^u y_{a,t} + B_1^u r_{a,t} + B_2^u w_t \\ \mathbf{y}_t^u &= C^u \mathbf{x}_t^u + \mathbf{r}_t^u\end{aligned} \tag{3.4}$$

since the measurement noise of the acceleration now must be taken into consideration as a process noise. The model can be rewritten as

$$\begin{aligned}\mathbf{x}_{t+T}^u &= A^u \mathbf{x}_t^u + B_1^u y_{a,t} + \begin{pmatrix} B_1^u & B_2^u \end{pmatrix} \begin{pmatrix} r_{a,t} \\ w_t \end{pmatrix} \\ \mathbf{y}_t^u &= C^u \mathbf{x}_t^u + \mathbf{r}_t^u .\end{aligned} \tag{3.5}$$

Since r_a and w are independent, the resulting process noise covariance is

$$Q^u = \begin{pmatrix} B_1^u & B_2^u \end{pmatrix} \begin{pmatrix} \sigma_a^2 & 0 \\ 0 & \sigma_w^2 \end{pmatrix} \begin{pmatrix} B_1^u \\ B_2^u \end{pmatrix}. \tag{3.6}$$

Since all measurements $\mathbf{y}^u$ in (3.5) should be deemed insignificant, the measurement covariances of $\mathbf{r}^u$ will be very large. To make the model comparison straightforward, the same noise parameters are used in both models. Therefore, the two noise parameters in (3.6) have not been replaced by a single parameter.

The reduced model becomes

$$\begin{aligned}\begin{pmatrix} x_{p,t+T} \\ x_{v,t+T} \end{pmatrix} &= \begin{pmatrix} 1 & T \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_{p,t} \\ x_{v,t} \end{pmatrix} + \begin{pmatrix} \frac{T^2}{2} \\ T \end{pmatrix} y_{a,t} + \begin{pmatrix} \frac{T^2}{2} & \frac{T^3}{6} \\ T & \frac{T^2}{2} \end{pmatrix} \begin{pmatrix} r_{a,t} \\ w_t \end{pmatrix} \\ \begin{pmatrix} y_{p,t} \\ y_{v,t} \end{pmatrix} &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x_{p,t} \\ x_{v,t} \end{pmatrix} + \begin{pmatrix} r_{p,t} \\ r_{v,t} \end{pmatrix}.\end{aligned} \tag{3.7}$$

In this model, the time update step will in practice provide all new information since σ_p^2 and σ_v^2 are very large.

3.2 Transfer Function Derivations

To compare the models, the transfer functions from the measurements to the estimates are derived. They are calculated using the stationary Kalman filter.

3.2.1 Transfer Function of Full Model

For the full model the transfer function is straightforwardly calculated using the time and measurement update equations of the Kalman filter.

The time update of the full state model is

$$\hat{\mathbf{x}}_{t+T|t} = A\hat{\mathbf{x}}_{t|t} \tag{3.8}$$

while the measurement update is

$$\hat{\mathbf{x}}_{t+T|t+T} = \hat{\mathbf{x}}_{t+T|t} + K_{t+T}\Big(\mathbf{y}_{t+T} - C\hat{\mathbf{x}}_{t+T|t}\Big). \tag{3.9}$$

For a stationary Kalman filter, the Kalman gain $K_{t+T} = \bar{K}$. $\bar{K}$ is calculated by first solving the Ricatti equation

$$\bar{P} = A\bar{P}A^T + Q - A\bar{P}C^T(C\bar{P}C^T + R)^{-1}C\bar{P}A^T \tag{3.10}$$

to calculate $\bar{P}$, and then determine $\bar{K}$ as

$$\bar{K} = \bar{P}C^T\Big(C\bar{P}C^T + R\Big)^{-1}. \tag{3.11}$$

The covariance matrices R and Q in (3.10) are

$$R = \begin{pmatrix} \sigma_p^2 & 0 & 0 \\ 0 & \sigma_v^2 & 0 \\ 0 & 0 & \sigma_a^2 \end{pmatrix}$$
$$Q = BB^T\sigma_w^2.$$

Unfortunately, there is no compact solution $\bar{P}$ to (3.10) as a function of T, σ_p^2, σ_v^2, σ_a^2 and σ_w^2, why $\bar{K}$ cannot be expressed as a function of these variables. Numerical solutions to (3.10) are though readily available for a given set of noise parameters and sampling time.

Introducing the q-operator, $qx(t) = x_{t+T}$ and $q^{-1}x(t) = x_{t-T}$, and using (3.8), (3.9) and $\bar{K}$, we can now derive the transfer functions $\mathbf{H}$.

$$\begin{aligned} \hat{\mathbf{x}}_{t+T|t+T} &= A\hat{\mathbf{x}}_{t|t} + \bar{K}\Big(\mathbf{y}_{t+T} - CA\hat{\mathbf{x}}_{t|t}\Big) && \Leftrightarrow \\ (qI - A + \bar{K}CA)\hat{\mathbf{x}}_{t|t} &= \bar{K}q\mathbf{y}_t && \Leftrightarrow \\ \hat{\mathbf{x}}_{t|t} &= \underbrace{(qI - A + \bar{K}CA)^{-1}\bar{K}q}_{\mathbf{H}(q)}\,\mathbf{y}_t && (3.12) \end{aligned}$$

where I is the identity matrix.

$\mathbf{H}(q)$ is the set of transfer functions $H_{ij}(q)$ which should be read as transfer func-

tions from measurement j to state i, where $i, j \in \{p, v, a\}$;

$$\begin{aligned}\hat{x}_p &= H_{pp}(q)y_p + H_{pv}(q)y_v + H_{pa}(q)y_a \\ \hat{x}_v &= H_{vp}(q)y_p + H_{vv}(q)y_v + H_{va}(q)y_a \\ \hat{x}_a &= H_{ap}(q)y_p + H_{av}(q)y_v + H_{aa}(q)y_a.\end{aligned} \tag{3.13}$$

The two most interesting transfer functions to study are $H_{pa}(q)$ and $H_{va}(q)$ since they reveal how the position and velocity estimates are depending on the acceleration measurements. They will be compared to the corresponding transfer functions of the input model.

3.2.2 Transfer Function of Reduced Model

The time update of the reduced model uses the acceleration measurement to estimate the future states.

$$\hat{\mathbf{x}}^u_{t+T|t} = A^u\hat{\mathbf{x}}_{t|t} + B^u_1 y_{a,t} \tag{3.14}$$

The measurement update uses the measurements $\mathbf{y}^u_{t+T}$ which the filter has absolutely no confidence in why they will only affect the state estimates in a minuscule way. Still, the update is needed to compute the transfer functions.

$$\hat{\mathbf{x}}^u_{t+T|t+T} = \hat{\mathbf{x}}^u_{t+T|t} + K^u_{t+T}\Big(\mathbf{y}^u_{t+T} - C^u\hat{\mathbf{x}}^u_{t+T|t}\Big). \tag{3.15}$$

Correspondingly to the full state model, $K^u_{t+T} = \bar{K}^u$ for a stationary filter and is calculated by using $\bar{P}^u$ that is acquired from the solution of the Ricatti equation as in Section 3.2.1. The necessary process noise Q^u was described in (3.6) and the measurement noise is

$$R^u = \begin{pmatrix}\sigma_p^2 & 0 \\ 0 & \sigma_v^2\end{pmatrix}. \tag{3.16}$$

Now, (3.14) inserted into (3.15), the q-operator and $\bar{K}^u$ gives us

$$\begin{aligned}\hat{\mathbf{x}}^u_{t+T|t+T} &= A^u\hat{\mathbf{x}}^u_{t|t} + B^u_1 y_{a,t} \\ &\quad + \bar{K}^u\Big(\mathbf{y}^u_{t+T} - C^uA^u\hat{\mathbf{x}}_{t|t} + C^uB^u_1 y_{a,t}\Big) && \Leftrightarrow \\ (qI - A^u + \bar{K}^uC^uA^u)\hat{\mathbf{x}}^u_{t|t} &= (B^u_1 + \bar{K}^uC^uB^u_1)y_{a,t} + \bar{K}^u q\mathbf{y}^u_t && \Leftrightarrow \\ \hat{\mathbf{x}}^u_{t|t} &= \underbrace{(qI - A^u + \bar{K}^uC^uA^u)^{-1}\bar{K}^u q}_{\mathbf{H}^u(q)}\,\mathbf{y}^u_t \\ &\quad + \underbrace{(qI - A^u + \bar{K}^uC^uA^u)^{-1}(I + \bar{K}^uC^u)B^u_1}_{\mathbf{H}^u_a(q)}\,y_{a,t}\end{aligned} \tag{3.17}$$

The transfer functions can be written as

$$\begin{aligned}\hat{x}_p &= H^u_{pp}(q)y_p + H^u_{pv}(q)y_v + H^u_{pa}(q)y_a \\ \hat{x}_v &= H^u_{vp}(q)y_p + H^u_{vv}(q)y_v + H^u_{va}(q)y_a\end{aligned} \tag{3.18}$$

where $H^u_{pa}(q)$ and $H^u_{va}(q)y_a$ stem from $\mathbf{H}^u_a(q)$ while the other four originate from $\mathbf{H}^u(q)$.

$H^u_{pa}(q)$ and $H^u_{va}(q)$ should be compared to $H_{pa}(q)$ and $H_{va}(q)$ to determine how the reduced model and the full state model differ. This will be done by studying the Bode plots of the transfer functions for a given parameters setting.

3.3 Bode Plots Evaluations

To evaluate the transfer functions the unknown parameters must be set. As explained in Section 3.1, σ_p^2 and σ_v^2 were chosen very large to ensure that their corresponding measurements do not affect the state estimates. They were set to $\sigma_p^2 = \sigma_v^2 = 10^{15}$. The sampling frequency was set to 100 Hz since it is common in for example IMUs. This gives $T = 0.01$.

Two parameters are now left; the process noise covariance σ_w^2 and the acceleration measurement noise covariance σ_a^2. As long as they are kept much smaller than σ_p^2 and σ_v^2, only the ratio between the two determines the filter estimates. One can therefore be fixed while the other is changed to see if and how the transfer functions change. Since the amplitude is of no importance in this case, the process noise covariance is fixed to $\sigma_w^2 = 1$ for convenience. Note that for the reduced system, the total process noise of the system is a linear combination of σ_w^2 and σ_a^2 why the ratio is much less important. Therefore, the transfer functions for the reduced system hardly change at all for different noise covariance ratios.

First, the measurement covariance noise is set to $\sigma_a^2 = 0.1$. The resulting Bode plots of the transfer functions from acceleration measurement to velocity estimate, H_{va} and H^u_{va}, are shown in Figure 3.1. The gain curve has the slope -1 for most frequencies which means a pure integration. For high frequencies the full model loses gain indicating that it does not trust inputs from these frequencies as much. The reduced model on the other hand integrates the acceleration measurements for all high frequencies. The gain loss for low frequencies experienced by the reduced model is related to the position and velocity measurements y_p and y_v and is therefore irrelevant since those measurements normally do not exist. The larger σ_p^2 and σ_v^2 are chosen, the lower the frequency where the input gain loss appears. Unfortunately, the Ricatti equation cannot be solved for arbitrarily large σ_p^2 and σ_v^2 in Matlab.

The transfer functions from acceleration to position estimate, H^u_{pa} and H_{pa}, are shown in Figure 3.2. The gain has a -2 slope indicating a double integration. Again, the full model has a gain loss for high frequency measurements since the model finds these less trustworthy while the reduced model double integrates almost all frequency components. The low frequency plateau in the reduced model is related to the position and velocity measurements.

The knee frequency, i.e. the frequency where the full model breaks off from its -1 slope in Figure 3.1 or its -2 slope in Figure 3.2, is related to the signal-to-noise ratio (SNR) σ_w^2/σ_a^2. A first indication of this is shown in Figure 3.3 where $\sigma_a^2 = 0.001$

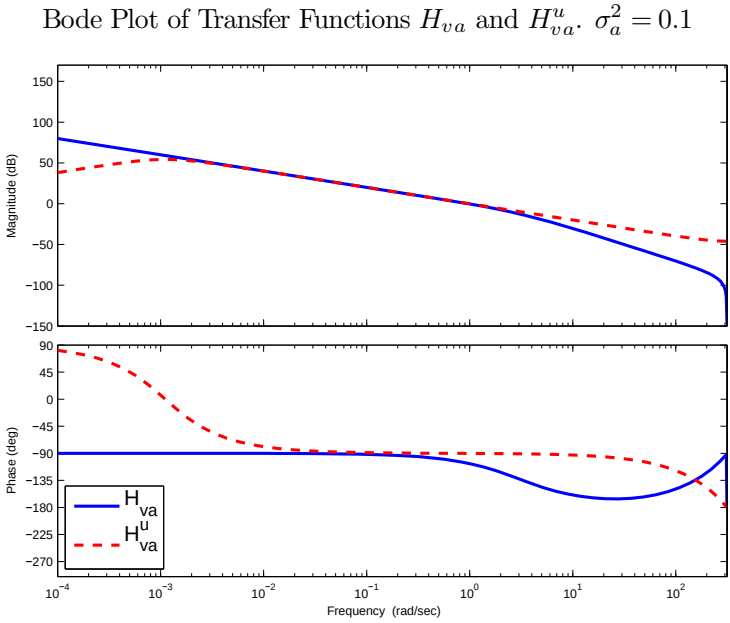

Figure 3.1: *Bode plot of transfer functions H^u_{va} and H_{va} for $\sigma_a^2 = 0.1$.*

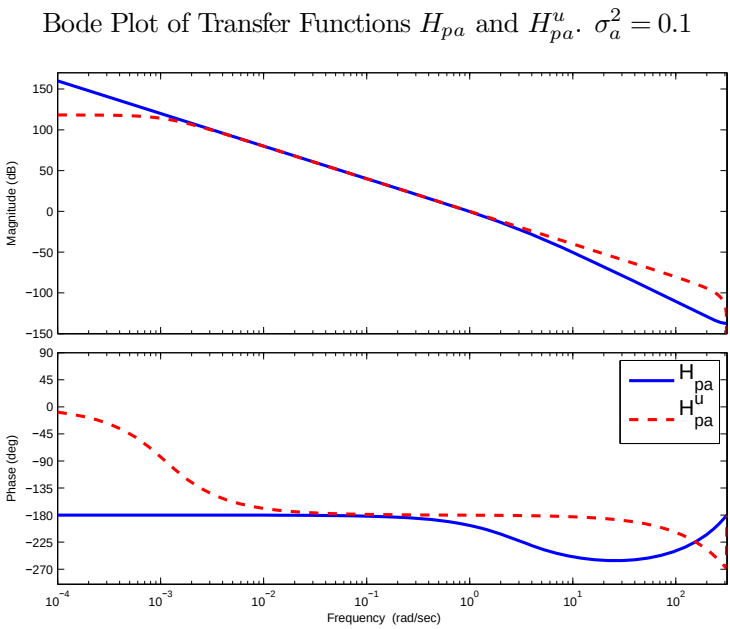

Figure 3.2: *Bode plot of transfer functions H^u_{pa} and H_{pa} for $\sigma_a^2 = 0.1$.*

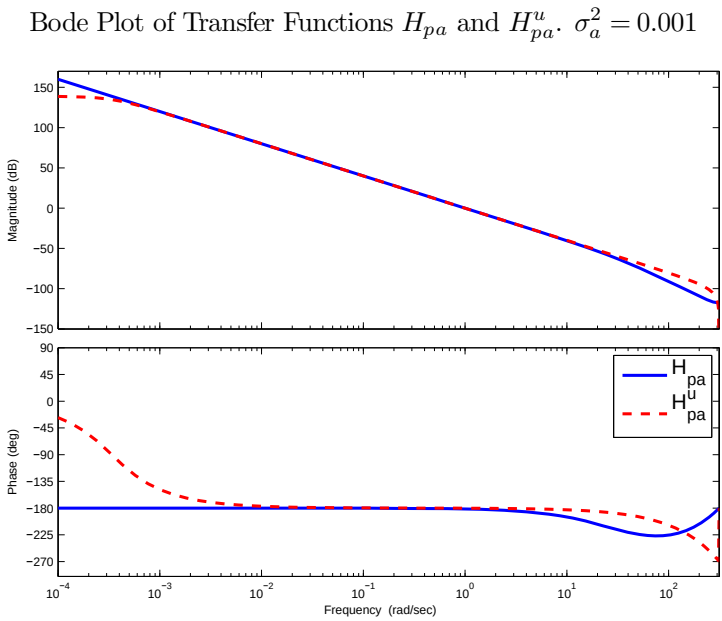

Figure 3.3: *Bode plot of transfer functions H^u_{pa} and H_{pa} for $\sigma_a^2 = 0.001$.*

giving an SNR a hundred times larger than in Figure 3.2. The knee frequency is here significantly higher than in Figure 3.2, roughly 35 rad/s compared to 4.5 rad/s before.

When the knee frequency is plotted against SNR, a clear connection is visible, see Figure 3.4. The relation between the 10-logarithm of the knee frequency β and the 10-logarithm of the SNR is almost perfectly linear. The exact frequency to choose as β is quite *ad hoc* though, especially for the low SNR cases. The linear function is estimated to be

$$\log_{10}(\beta) = 0.5 \log_{10}(\text{SNR}) + 0.05. \tag{3.19}$$

This knee frequency is related to the positions of the system's poles. There are always two poles very close to 1 resulting in a double integration. A third pole is found on the real axis in the unit circle. For lower SNR, the third pole wanders closer to 1 causing the gain curve to drop for lower frequencies. The pole placement of H_{pa} is shown using a root locus in Figure 3.5.

3.4 Discussion

The main drawback of using the reduced model instead of the full model is that the position and velocity estimates become more sensitive to high frequency noise in the acceleration measurement. Everything is just accepted as true user accelerations and is integrated straight into the states. The full model has an additional

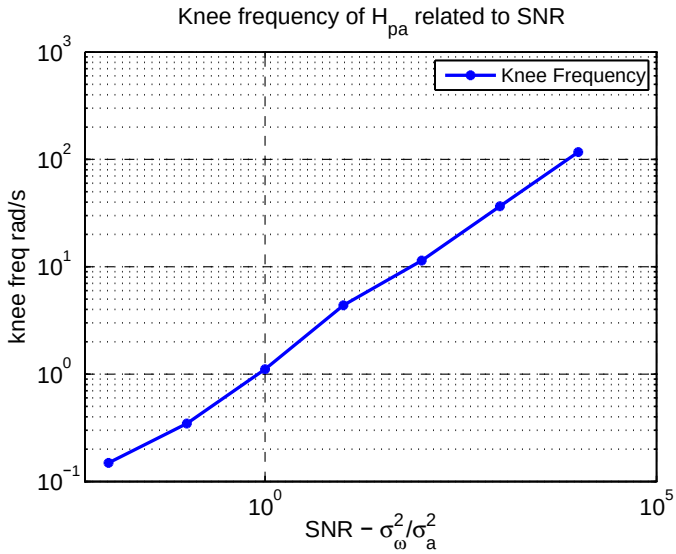

Figure 3.4: *The knee frequency of H_{pa} and H_{va} related to SNR.*

step that prefilters the acceleration measurement before it enters the state estimates making them less noisy. The full filter estimates are therefore a tradeoff between the trustworthiness of the model and the measurements, while the reduced model has no such tweaking abilities. The only thing one can influence in the reduced model are the sizes of the state covariances, not the estimates themselves.

One might get the idea of low pass filtering the acceleration measurements before they are fed to the reduced model to compensate for the high frequency gain difference. Unfortunately this will only result in the measurement noise r_a being frequency dependent. This in turn would then have to be compensated for in the reduced filter why the low pass filter will only solve the problem by introducing a new one.

Another option is to use a different sampling technique like *first-order hold* or the *bilinear transformation*. This might smooth out the acceleration measurements and thereby dampen the high frequency effects of the *zero-order hold* sampling that is used.

Since the main advantage of the reduced model is its simpler implementation, the application must specify which model structure to use. If computational resources are scarce, the reduced model might be preferable despite its larger sensitivity, especially if the reduced structure can be applied in multiple dimensions.

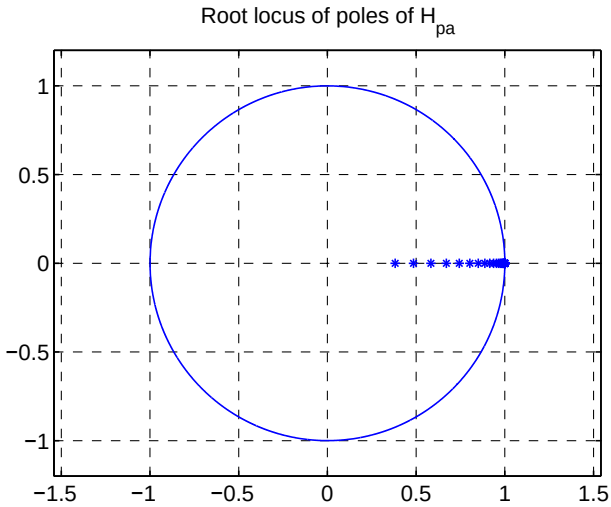

Figure 3.5: *Root locus of the poles of H_{pa} for different* SNR*. For lower* SNR*, the pole on the real axis moves closer to* 1*. Two poles are always very close to* 1*.*

4

Indoor Localization

There is a growing desire in some fields to track people moving around in a building. When firefighters, police officers or soldiers search through a building, accurately knowing the true position of the personnel would greatly enhance their safety. Finding a trapped firefighter in a burning building is far easier if his or her position is available for the rescue team, with or without a map.

The problem of indoor localization has received an increasing amount of attention in the last couple of years; Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007); Godha et al. (2006); Woodman and Harle (2009); Grzonka et al. (2010); Aggarwal et al. (2011); Jiménez et al. (2010); Robertson et al. (2009); Widyawan et al. (2008). Since GPS signals cannot be used indoors due to the extreme weakness of the signals, new positioning techniques must be developed. The drive is to put sensors in the first responders' uniforms and use these to track the movements and thereby the position of the user. This has led to a transition from robot sensor platforms to human ones.

In robotics, localization spans all types of robots which can be equipped with a wide variety of sensors and operate in a multitude of environments. Particularly the SLAM problem, estimating a map while positioning oneself in the same map, has been well studied. The land based robot commonly used for this type of problem is equipped with wheel encoders, laser sensors, cameras, accelerometers, etc. A robot is a neat platform since the structural relationships between the sensors are constant and the sensors are almost always observing the surroundings from the same height and tilt.

The human sensor platform provides new challenges due to the non-sturdiness of the human body compared to a robot. If the person is equipped with multiple sensors on different positions on the body, the distances between these sensors

will not be fixed. Thus, the mutual positions of all sensors can vary, but only in restricted ways. For example, the distance from the head to the feet can be shorter than the full length of the user if he or she is crouching, but there are restrictions on how tall an uninjured user can be.

To track a person, a variety of sensors can be used. An IMU with accelerometers and gyros and maybe also magnetometers is simple and cheap and is therefore a very common sensor. Unfortunately, a positioning system that only uses accelerometers and gyros is prone to be quite wrong quite soon why it is usually supported by a range measuring radio device such as WiFi or Ultra Wide Band, Woodman and Harle (2009), or is fused with preexisting maps for enhanced tracking precision, Woodman and Harle (2009); Aggarwal et al. (2011); Widyawan et al. (2008).

The IMUs used for indoor localization are small and cheap and consequently perform quite poorly. There is commonly a drift in the gyros causing the orientation estimate to be incorrect after a short period of time. Since the orientation is wrong, the direction of down is wrong and a part of the gravitational acceleration will instead be believed to originate from the user movements. This error is double integrated to estimate the sensor position, resulting in a positioning error that grows cubically in time. This rapidly causes very large positioning errors.

By mounting the IMU on the foot, the positioning errors can be greatly reduced. If we can safely detect that the foot is on the ground we know that it is at a stand still which we can use in the filter. Since the foot is stationary during the stance, virtual measurements of zero velocity can be fed to the positioning framework. This is known as zero velocity update (ZUPT) and reduces the positioning error from being cubical to linear in time Foxlin (2005).

The problem of false positives, i.e. calling a stand still when the foot is in midair, is a serious one. If the foot is not at a standstill when one is called, the filter will be updated with false virtual measurements which can ruin the position estimates. False negatives, i.e. missing one or two true stand stills, is not an equally big problem. The stand still detection framework should therefore be designed with the focus on minimizing the false positives.

This chapter consists of three parts. In the first section, Section 4.1, a probabilistic framework for stand still detection is derived. The second section, Section 4.2, covers an extensive investigation of how the position of the IMU on the boot affects the possibility to detect stand stills for a number of user movements. In the final section, Section 4.3, the stand still detection framework is incorporated in a localization framework that uses ZUPT to position a user. The positioning experiments were performed in the VICON lab that provides accurate ground truth of the movements.

4.1 Stand Still Detection

The key to making a positioning system with a foot mounted IMU work is to safely detect that the foot is at a stand still. True detected stand stills used for ZUPT will greatly enhance the localization performance. False stand stills used for ZUPT will though risk causing large positioning errors. It is therefore worse to detect false positives than to miss true positives. In this section we will look at a probabilistic framework for stand still detection.

To detect that the foot is touching the ground, one obvious option would be to use pressure sensors mounted under the soles. There are though some drawbacks with this solution. First, pressure sensors usually have some sort of moving parts which are prone to be worn out by usage. Second, there are cases when the foot is at a standstill even though the sole is not touching the ground, for example when the subject is crawling, why the stationary phases would not be detected by the pressure sensors. And third, if the foot is turning while on the ground it would be regarded as a stand still by the pressure sensors even though it is moving. The drive is therefore to use the data from the foot mounted IMU itself to detect that the sensor is at a standstill.

Previously, stand still detection has mostly been done by comparing the accelerometer or gyro signal to a threshold. Here we put the stand still detection in a probabilistic framework using test statistics with known distributions and a Hidden Markov Model. The result is an estimated probability of stand still that is consecutively calculated at every time instant. This can be used for ZUPT in a filtering framework. This section is based on the work presented in Callmer et al. (2010b) and Rantakokko et al. (2011).

4.1.1 Related Work

Most solutions to the stand still detection problem use an averaged accelerometer or gyro measurement and compares it to a threshold; Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007). The threshold is chosen *ad hoc* and is normally quite restrictive to minimize false positives. Another approach is the moving variance used in Godha et al. (2006) where the variance computed over a sliding window is compared to a threshold. One of the problems with this approach is that in order to make it work properly the time interval during which the boot is still must be quite long. Therefore it might work well detecting stand stills during walking but not during running.

Probabilistic zero velocity detection has previously been proposed in Skog (2009) who used a hypothesis test to determine if the foot was stationary or moving. The hypothesis test was performed using a test statistic based on a Generalized Likelihood Ratio Test (GLRT). The pdf of the acceleration and/or the angular velocity during the swing phase of the step, was approximated with an unnormalized uniform distribution. The pdf during stand still was based on the exponential of the norm of the acceleration and/or the angular velocity, which has an unknown distribution. The resulting test statistic was a moving average of the norm of the

acceleration measurements and/or the angular velocity measurements. This was compared to a threshold to determine if the foot was to be rendered stationary. Since the test statistic has an unknown distribution the threshold was chosen *ad hoc*, making the framework in practice similar to the ones in Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007).

The test statistics used in Skog (2009) are similar to the ones used in this work since we both evaluate one acceleration based and one angular velocity based. The acceleration based test statistic differs though in that we have chosen one which has a known distribution. Skog (2009) also has a third test statistic based on a combination of acceleration and angular velocity while we choose to combine the measurements by first evaluating the signals separately and then fusing the estimated probabilities.

4.1.2 Test Statistics Derivation

The test statistics are based on the accelerometer and gyro signals from the IMU. The sensor was mounted on the boot under the shoe laces. An example of a measured walking sequence with a shoe lace mounted IMU is shown in Figure 4.1. The foot is stationary around time instants 550, 660, 770, 870 and 980. During these phases the norm of the accelerometer signals is the gravitation constant with some added noise. At the same time instants the norm of the angular velocity signal is zero with some additive noise. These will be the key characteristics of the test statistics.

Sensor Models

The signal model is

$$\mathbf{y}_t = \begin{bmatrix} \mathbf{y}_t^a(\theta) \\ \mathbf{y}_t^\omega(\theta) \end{bmatrix} + \begin{bmatrix} \mathbf{v}_t^a \\ \mathbf{v}_t^\omega \end{bmatrix} \tag{4.1}$$

where $\mathbf{y}_t^a$ and $\mathbf{y}_t^\omega$ denote the acceleration measurement vector and the angular velocity measurement vector, respectively. Further, θ denotes the model dependence on the phase of the human step sequence. Naturally, the model differs significantly between when the foot is at stand still and when it is moving.

The measurements are assumed to have additive independent identically distributed Gaussian noise $\mathbf{v}^a \sim N(\mathbf{0}, \boldsymbol{\sigma}_a^2)$ and $\mathbf{v}^\omega \sim N(\mathbf{0}, \boldsymbol{\sigma}_\omega^2)$ where $\boldsymbol{\sigma}_\omega^2 = \sigma_\omega^2 \mathbf{I}$ and $\boldsymbol{\sigma}_a^2 = \sigma_a^2 \mathbf{I}$ and $\mathbf{I}$ is the 3×3 identity matrix.

During stand still the sensor model is

$$\begin{bmatrix} \mathbf{y}_t^a \\ \mathbf{y}_t^\omega \end{bmatrix} = \begin{bmatrix} g\mathbf{u}_t \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{v}_t^a \\ \mathbf{v}_t^\omega \end{bmatrix}, \tag{4.2}$$

where $\mathbf{u}_t$ is the unknown gravitational direction vector and g is the gravitational constant 9.81. Since the orientation of the boot changes over time, so does $\mathbf{u}_t$. When the foot is moving the sensor model changes to

$$\begin{bmatrix} \mathbf{y}_t^a \\ \mathbf{y}_t^\omega \end{bmatrix} = \begin{bmatrix} g\mathbf{u}_t + \mathbf{a}_t \\ \boldsymbol{\omega}_t \end{bmatrix} + \begin{bmatrix} \mathbf{v}_t^a \\ \mathbf{v}_t^\omega \end{bmatrix} \tag{4.3}$$

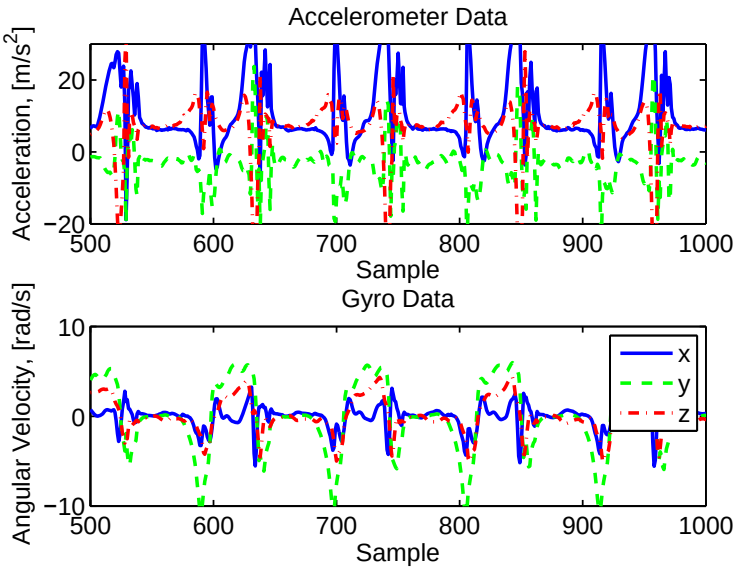

Figure 4.1: *Example of accelerometer data (where x, y and z is solid, dashed and dashdotted) and gyro data (ω_x, ω_y and ω_z is solid, dashed and dash-dotted, ω_i is angular rotation rate around axis i) during a walking sequence. The foot is stationary around time instants 550, 660, 770, 870 and 980.*

where $\mathbf{a}_t$ and $\boldsymbol{\omega}_t$ are induced by human movements and therefore have unknown distributions.

All stand still detection frameworks are roughly based on the assumption that $\mathbf{a}_t$ and $\boldsymbol{\omega}_t$ are not both zero at the same time instances if the foot is not stationary.

Test Statistics

To be able to differentiate between the two modes, test statistics with known distributions are computed. Two different ones are evaluated here, one using only the accelerometer data, T^a, and one using only the angular velocity data, T^ω.

The test statistic of the accelerometer magnitude detector is computed as

$$T_t^a = \frac{\|\mathbf{y}_t^a\|^2}{\sigma_a^2} \tag{4.4}$$

where $T^a \sim \chi^2(3, \lambda)$ during stand still. It has a noncentral chi-square distribution since $\mathbf{y}_t^a$ has nonzero mean when the foot is stationary. Its noncentrality parameter $\lambda = g^2/\sigma_a^2$ and 3 is the number of degrees of freedom.

The angular velocity test statistic is

$$T_t^{\omega} = \frac{\|\mathbf{y}_t^{\omega}\|^2}{\sigma_{\omega}^2} \tag{4.5}$$

where $T^{\omega} \sim \chi^2(3)$ during stand still since $\mathbf{y}^{\omega}$ has zero mean when the foot is stationary. T^{ω} has 3 degrees of freedom.

Test Statistic Appearance during Walking Sequence

The two test statistics are computed for the walking sequence in Figure 4.1 and is shown plotted in logscale in Figure 4.2. A plot of the test statistics of the walking sequence in The means of the stand still distributions are marked with dashed lines. The stand still events occurring around time instants 550, 660, 770, 870 and 980 are clearly visible. The test statistic T^a has a movement distribution

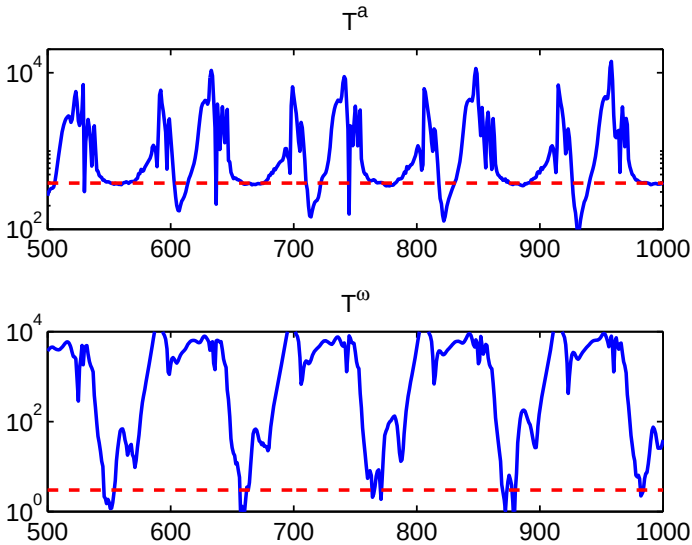

Figure 4.2: *Logarithmic plot of the test statistics with the mean of the stand still distribution marked with a dashed line. T_t^a at the top and T_t^{ω} at the bottom. The foot is stationary around time instants 550, 660, 770, 870 and 980.*

that has a significant overlap with the stand still distribution, causing the test statistic to cross the mean of the stand still distribution during the stride. This is shown around time instants 530, 615, 630, 710, 750, 825 and 935. Simply calling a stand still when T^a is close to the mean of the stand still distribution will therefore cause a lot of false positives.

T^{ω} has a distribution during movement that does not have a real overlap of the stand still distribution, why T^{ω} may be a safer test statistic than T^a to use for

stand still detection. Figure 4.2 also shows that the stand still phases of T^{ω} appear shorter than indicated by T^{a}. Apparently, the angular velocity changes before the acceleration does when a step is started.

4.1.3 Test Statistic Distribution Validation

The test statistics must be validated to ensure that their experimental stand still distributions are similar to the theoretical ones. We also estimate the distributions of the test statistics under experimental movements and plot how they relate to the stand still distributions. The empirical movement distributions are shown as histograms and are plotted with their respective approximations. The histograms were created using a large amount of experimental data. The same movement distributions were used for all types of movements since the differences between them were quite small, especially in the most important regions close to the stand still distributions.

Acceleration Magnitude Detector

The distributions of the acceleration magnitude test statistic T^{a} are shown in Figure 4.3. The theoretical and the empirical stand still distributions have similar mean but slightly different covariances. One of the reasons why the empirical density has a smaller covariance than the theoretical one, could be that we have been a bit too meticulous selecting the stand still data. Note also the significant overlap of the probability distributions of stand still and movement. That makes it difficult to safely identify stand stills by only looking at T^{a}.

The movement distribution was approximated using two Gaussians as

$$p_m(T^{a}) = 0.94 \cdot N(500, 1000) + 0.06 \cdot N(390, 10).$$

Angular Rate Magnitude Detector

The distributions of T^{ω} are shown in Figure 4.4. Clearly, the theoretical stand still distribution is very similar to the empirical one, estimated by experimental data. The movement distribution was approximated using two Gaussians as

$$p_m(T^{\omega}) = 0.25 \cdot N(30, 25) + N(5000, 3000).$$

Also note the much smaller overlap between the stand still and movement distributions, bottom left in Figure 4.4. This should enable more robust stand still detection when using T^{ω} instead of T^{a}.

4.1.4 Hidden Markov Model

To determine the probability of stand still, a Hidden Markov Model (HMM) is used. It is based on modes, in this case two modes; stand still and movement, and it estimates the probability of being in each mode at every time instant in a recursive fashion. To do that it uses a predefined probability of switching between the modes, the distributions of the test statistic under each mode and the measured value of the test statistic.

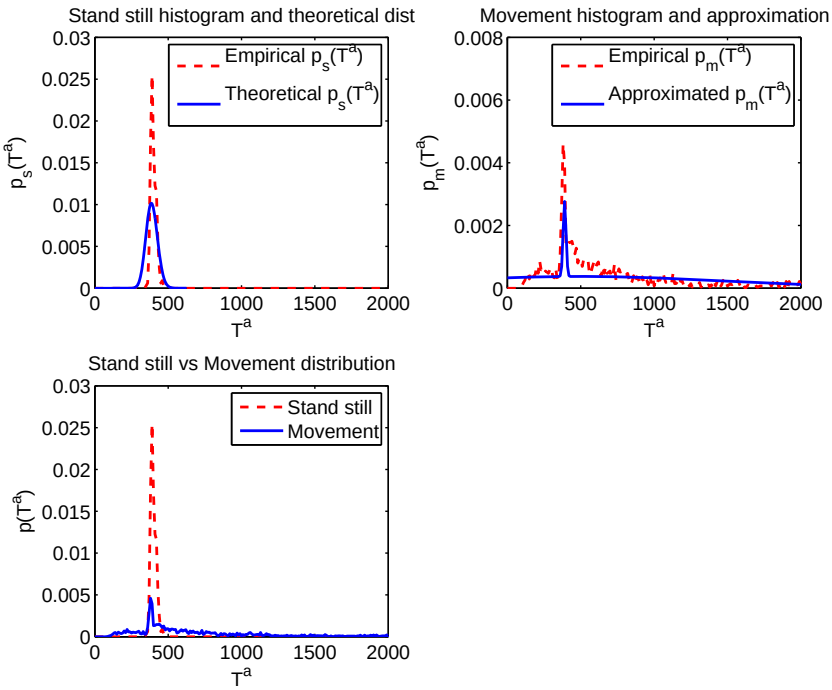

Figure 4.3: *Top left: theoretical stand still distribution of T^a (solid) and empirical histogram (dashed), top right: empirical movement histogram (dashed) and approximation (solid), bottom left: empirical stand still (dashed) and movement (solid) distributions. Note the significant overlap between the stand still and movement distribution.*

One question might arise at this point. Why cannot we just formulate a simple hypothesis test to figure out if the foot is stationary? First of all, a hypothesis test only provides some sort of information if the H_0 hypothesis is rejected. Otherwise we do not know anything. That means that in order to detect stand stills, the H_0 hypothesis must be that the foot is moving. This in turn means that we must reject the H_0 hypothesis based on the test statistic being below or above a certain threshold. This threshold is based on the movement distribution which is unknown. Since we cannot determine a statistically derived threshold, a hypothesis test cannot be formulated.

The HMM has as stated above two modes; mode 1 when the foot is at a stand still and mode 2 when the foot is moving. The mode transition probability matrix states the probability of a mode switch which induces some dynamics into the probability estimation. A lower mode transition probability requires a measurement with a higher likelihood for a mode switch to occur.

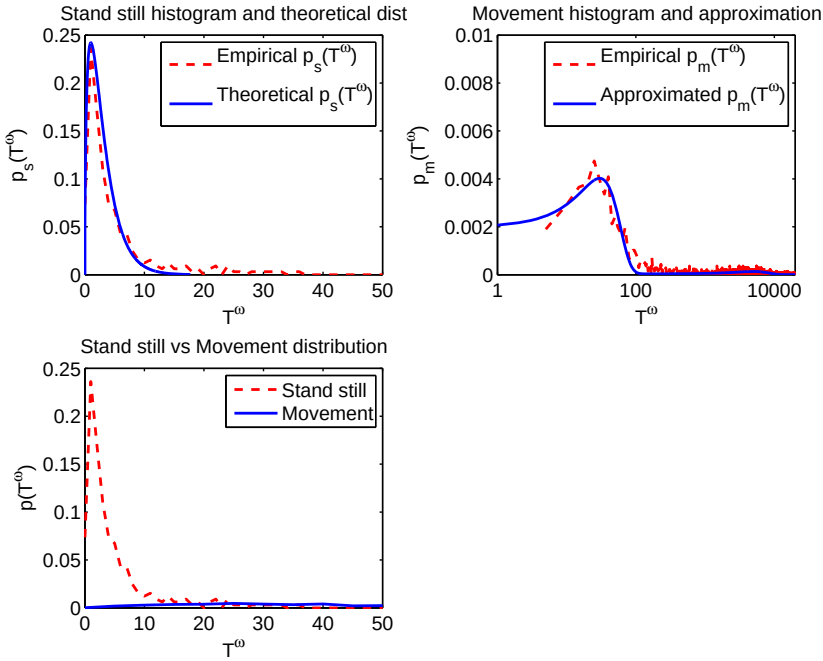

Figure 4.4: *Top left: theoretical stand still distribution of T^ω (solid) and empirical histogram (dashed), top right: empirical movement histogram (dashed) and approximation (solid), bottom left: empirical stand still (dashed) and movement (solid) distributions. Note the difference in magnitude between the stand still and movement distribution (bottom left).*

The mode transition probability matrix used in the experiment is

$$\Pi = \begin{bmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{bmatrix} \tag{4.6}$$

which states that the probability of going from stand still to moving or vice versa, is 5%. During normal walking the right foot takes about one step per second which results in roughly 2 mode transitions every 100 measurements. The transition probabilities were chosen slightly higher to incorporate also faster movements.

The mode probabilities at time t are calculated using the recursion

$$\begin{aligned} \mu_t^i &= P(r_t = i|\mathbf{y}_t) \propto p(\mathbf{y}_t|r_t = i)P(r_t = i|\mathbf{y}_{t-1}) \\ &= p(T_t|r_t = i) \sum_{j=1}^{N_r} \Pi_{ji}\mu_{t-1}^j. \end{aligned} \tag{4.7}$$

where r_t is the mode state. Hence we have

$$\mu_t^i = \frac{p(T_t|r_t = i)\sum_{j=1}^{N_r} \Pi_{ji}\mu_{t-1}^j}{\sum_{l=1}^{N_r} p(T_t|r_t = l)\sum_{j=1}^{N_r} \Pi_{jl}\mu_{t-1}^j}. \tag{4.8}$$

The probability density functions of movement used in the HMM are approximations set to resemble the empirical movement density functions in Figures 4.3 and 4.4. The HMM framework thus consecutively estimates the probability of movement and stand still for each time instant.

If the probability transition matrix is chosen as having the same probability of change between all modes, i.e.

$$\Pi^c = \begin{bmatrix} \alpha & \dots & \alpha \\ \vdots & \ddots & \vdots \\ \alpha & \dots & \alpha \end{bmatrix} \tag{4.9}$$

the probability of each mode is only depending on the latest measurement.

$$\begin{aligned} \mu_t^i &= \frac{p(T_t|r_t = i)\sum_{j=1}^{N_r} \Pi_{ji}^c\mu_{t-1}^j}{\sum_{l=1}^{N_r} p(T_t|r_t = l)\sum_{j=1}^{N_r} \Pi_{jl}^c\mu_{t-1}^j} \\ &= \frac{p(T_t|r_t = i)\alpha\sum_{j=1}^{N_r} \mu_{t-1}^j}{\sum_{l=1}^{N_r} p(T_t|r_t = l)\alpha\sum_{j=1}^{N_r} \mu_{t-1}^j} \\ &= \frac{p(T_t|r_t = i)}{\sum_{l=1}^{N_r} p(T_t|r_t = l)} \end{aligned} \tag{4.10}$$

All dynamics are then removed making it prone to give false positives by only one troublesome measurement. This is the same as saying that a stand still can occur at any time and that each measurement is independent of all previous measurements which is not true.

For the two mode case stand still and movement, the two step algorithm can now be summarized in Algorithm 3.

Combined Test Statistic

We now have one mode probability based on the accelerometer test statistic and one on the gyro test statistic. Instead of just looking at one of these, one can fuse these mode probabilities to get a stand still probability based on both acceleration and angular velocity. Theoretically, this would reduce the number of false positives since there is a small chance that both will be wrong at the same time. On the other hand, if one of the test statistics fails to indicate the true positives for some reason, all true positives will be missed.

Algorithm 3 Stand Still Detection Framework

Require: Stand still distribution p_s and movement distribution p_m of test statistic T. Measurements $\mathbf{y}$ with noise parameter σ. Mode transition probability matrix Π.

1: **for** $t = t_{\text{start}}, \ldots, t_{\text{final}}$ **do**

2: Compute test statistic.

$$T_t = \frac{\|\mathbf{y}_t\|^2}{\sigma^2} \tag{4.11}$$

3: Estimate stand still and movement probabilities, μ_t^s and μ_t^m.

$$\mu_t^s = \frac{P_{s,t}}{P_{s,t} + P_{m,t}} \qquad \mu_t^m = \frac{P_{m,t}}{P_{s,t} + P_{m,t}}$$

where

$$\begin{aligned} P_{s,t} &= p_s(T_t)\left(\Pi_{ss}\mu_{t-1}^s + \Pi_{sm}\mu_{t-1}^m\right) \\ P_{m,t} &= p_m(T_t)\left(\Pi_{ms}\mu_{t-1}^s + \Pi_{mm}\mu_{t-1}^m\right) \end{aligned} \tag{4.12}$$

4: Declare stand still if $\mu_t^s \geq \alpha$.

5: **end for**

The combined stand still probabilities are calculated as

$$\mu_t^s = \mu_t^s(T^a) \cdot \mu_t^s(T^\omega) \tag{4.13}$$

making the movement probabilities

$$\mu_t^m = 1 - \mu_t^s. \tag{4.14}$$

In the experiments, this third option will be evaluated alongside the first two.

4.1.5 Experimental Results

The mode probabilities provided by the HMM of the data sequence in Figure 4.1 is shown in Figure 4.5. The top two show the probabilities indicated by T^a and T^ω and illustrate the difference in stand still detection performance. The bottom one is the combined one that only indicates a stand still when both test statistics indicate that there is a stand still.

In the experiments, α in Algorithm 3 was chosen as $\alpha = 0.99$. That means that we defined false positive as when the system calculated a stand still probability of at least a 99% while the foot was not at rest. If the foot actually was at rest the stand still was defined as detected. The measurement noises were set as $\sigma_a = 0.5$ and $\sigma_\omega = 0.08$, estimated from data when the sensor was at rest.

It is clear that the acceleration based test statistic T^a is close to indicating false positives around some of the troublesome time instants mentioned in Section 4.1.2; 500, 615 and 825. The framework does not call it a stand still after the first measurement of T^a that is close to the stand still mean, but after a couple of close measurements the HMM starts to assume that the foot is at rest.

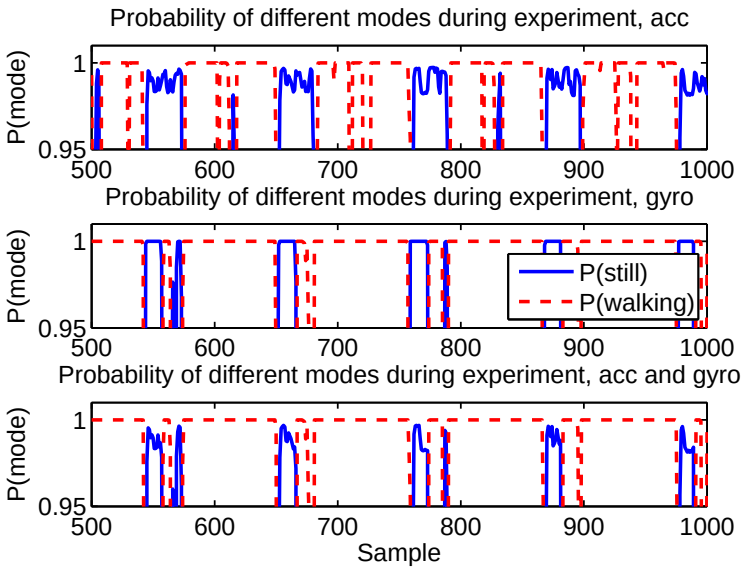

Figure 4.5: *From top to bottom, mode probabilities for T^a, T^ω and the combination, evaluated on the data set in Figure 4.1. The foot is stationary around time instants 550, 660, 770, 870 and 980.*

The angular velocity based test statistic T^ω gives a distinct detection of every foot stance. The stationary moment is rather short but is often followed by a shorter second stationary moment. Figure 4.2 shows that this is because there is commonly a slight angular movement halfway through the deemed stationary part. This second stationary moment provides no new information and only the first detection is necessary to perform ZUPT. No false positives occur during the stride phase of the step.

When both T^a and T^ω has to indicate that the foot is at a stand still, there are fewer detected stand stills. In Figure 4.2, we can see that it pretty much gives the same result as T^ω does. The combination does not seem necessary when walking, but it might be necessary if certain movements result in false positives using only the gyro.

Further walking experiments reveal the stand still detection performance shown in Table 4.1. All 190 true stationary phases were detected, but also some false positives.

Clearly, the acceleration based test statistic performs the worst and the combination is the best. The results are though only approximate. One of the big problems when evaluating indoor navigation systems is how to acquire reliable ground truth data. You can count the number of steps you take but that only works when you are just walking straight ahead and that is the least interesting

case since it is the simplest. A data set consisting of 100 steps when you close and open doors, move around furniture and such is much more difficult to analyze afterwards and it is also much more prone to result in false positives. This problem comes back over and over again when dealing with indoor navigation systems; what is the ground truth?

Walking, 190 steps	T^a	T^ω	both
Stand stills detected	190	190	190
False positives	128	12	5

Table 4.1: *True positives and false positives while detecting stand still phases. 190 steps were taken in 4 different experiments.*

4.1.6 Conclusions and Future Work

Two test statistics with known distributions have been evaluated for stand still detection. In conjunction with a Hidden Markov Model, the mode probabilities are readily calculated and can be used for zero velocity updates.

The test statistic based on accelerometer data only has been shown to result in plenty of false positives. This is natural since there is a significant overlap between the test statistic pdf during stand still and the pdf during movements. The gyro based one has been shown to provide good stand still detection capabilities during walking even though some false positives do occur, mainly when small movements are made. Combining the estimated stand still probabilities seems to be the safest stand still detector, since sometimes the foot may look like it is still when only one signal is considered but when looking at both signals you can tell that it is not.

One can also consider extending the test statistics by incorporating a number of measurements over time.

$$T_t^\omega = \sum_{i=t-k}^{t} \frac{\|\mathbf{y}_i^\omega\|^2}{\sigma_\omega^2} \tag{4.15}$$

During stand still the measurements will be independent. A short time interval, $k \approx 3$, might remove false positives resulting from just one or two troublesome measurements.

4.2 Stand Still Detection Performance for Different IMU Positions

For a localization framework with a foot mounted IMU, the sensor must be mounted on the shoe. Where to attach the IMU on the boot might seem like a straightforward problem; put it on the shoe of course! It is though not that straightforward. Different sections of the shoe are stationary during different parts of the stance. Also, different movements cause different sections of the boot to be stationary. A placement that is good for walking is not necessarily good for running or crawling. More motions than walking have to be taken into consideration since the end users of a professional indoor positioning system, i.e. firefighters, police officers and soldiers among others, have a tendency to do a lot more than walking. In this section we will investigate how the sensor position and shoe type affect the stand still detection performance using four different movements.

The first sensor position is by the heel, Figure 4.6a. In a real end user application one idea is to hide the sensor in the thick sole of the heel to protect it. Another option is to put the sensor by the toes since they are involved in almost all movements, Figure 4.6b. Intuitively, the toes should be stationary during pretty much all conceivable standing motions sequences. Since the sensors are getting smaller and smaller, the problem of hiding the sensor in the thinner sole by the toe could soon be solved. A third option is to put the sensor by the shoelaces, Figure 4.6c. This position could be common if the sensor is not integrated with the shoe, but is instead strapped onto the preexisting boot.

Two different types of boots have also been used; one softer winter boot, Figure 4.6b, and one rigid hiking boot, Figure 4.6a. The reason is to illustrate not only the importance of the sensor position, but also of the softness of the shoe. A softer shoe tends to be reshaped more by the ground during the stance making a larger section of the boot stationary at a given stance. A more rigid boot does not get reshaped during the stance making the boot more likely to roll against the ground. Therefore a rigid boot is often not really stationary during a stance. The likely end users of a foot mounted positioning system probably have boots with a rigidness somewhere in between the winter and the hiking boot. All three sensor positions will be evaluated on both boots using walking, running, duck walking and duck-drag. The last movement is common among firefighters and means that he or she is moving standing on one knee and one foot. The pose is shown by firefighter 55 in Figure 4.7. The sensor was assumed attached to the boot that is touching the ground with the sole. Duck walking is walking in a squatted position.

The main problem with evaluating even more realistic firefighter movements like crawling is that it is very difficult to obtain reliable ground truth data for such experiments. The experiment has to take place in for example a VICON lab, Section 2.1.4, but not even these experiments will be completely like crawling around furniture and such in an apartment. The data is somewhat periodic but still the stand still phases are not straightforward to pinpoint. The more distinct

(a) Hiking boot with the sensor by the heel. *(b) Winter boot with the sensor by the toe.* *(c) Hiking boot with the sensor by the shoe laces.*

Figure 4.6: *Sensor positions and boots used during stand still detection experiments. All combinations were evaluated.*

and repetitive the movement is, the easier it is to differentiate between true and false positives using nothing but the sensor data. Walking and also running experiments are therefore easy to evaluate, while crawling or duck-drag are more difficult, especially the former.

For all sensor positions and movements we will also compare the stand still detection performance of the gyro based test statistic and the combined one.

4.2.1 Experimental Results

All results are plotted as detection rate versus false positives rate in Figure 4.8. The top row shows walking experiments, second is running, third is duck walking and the last one is duck-dragging. The left column are the results using the softer winter boot and the right column are the results using the hiking boot. The stand still detection results using the angular velocity test statistic and the results

Figure 4.7: *Firefighter 55 is demonstrating the duck-drag position.*

using the combination are shown in the same plot. The same symbol is used for the same movement but the test statistics are color coded.

A stand still is assumed detected if the estimated stand still probability is above 99% during a true stand still. If no stand still did occur it is considered a false positive. Figure 4.8 is therefore a ROC-curve with only one threshold sample (99%).

The best result is a combination of high stand still detection ability combined with few false positives. This result has been achieved when the marker is plotted in the top left corner. The further to the right, the more false detections. The further down, the fewer stand still detections. All in all, the best combination seems to be a toe mounted sensor with a test statistic using a combination of angular velocity and acceleration measurements.

The complete results are also shown in table form in the end of the chapter. In each table the results using the softer winter boot are shown at the top and at the bottom are the results from the hiking boot experiments.

Walking

The first movement that will be evaluated is walking. This is the most common movement used in localization experiments, Beauregard (2007); Feliz et al. (2009); Foxlin (2005); Ojeda and Borenstein (2007); Godha et al. (2006); Skog (2009); Woodman and Harle (2009). This is probably because walking is easy to perform, the data is straightforward to analyze and the stand still phases are long and stable making them easy to detect.

The results from the walking sessions are shown in the top row in Figure 4.8 and in Table 4.2. It is clear that almost all stand still phases are detected no matter the boot or the sensor position. What is not shown in in Figure 4.8 and Table 4.2 are the characteristics of the stand still phases. With a heel mounted sensor the stand still phases were always quite short making them more difficult to detect. For a toe mounted sensor the stand still phases were long but sometimes the toes can have zero angular velocity when the boot is in mid air causing false positives. It might therefore be necessary to support the gyro based test statistic with the accelerometer based one to achieve truly safe stand still detection.

Running

The second movement is running which is a bit more difficult to analyze. It is periodical like walking which is nice, but the stand still sequences are significantly shorter making the true stand still moments more difficult to distinguish. A higher sampling frequency could help but it is not necessary that it would solve all the problems.

The second row in Figure 4.8 and Table 4.3 shows the stand still detection performance during running which shows clear differences between the settings. The stand still phases are pretty much non existent when a rigid hiking boot is used. If only accelerometer data was considered, some stand stills could still have been

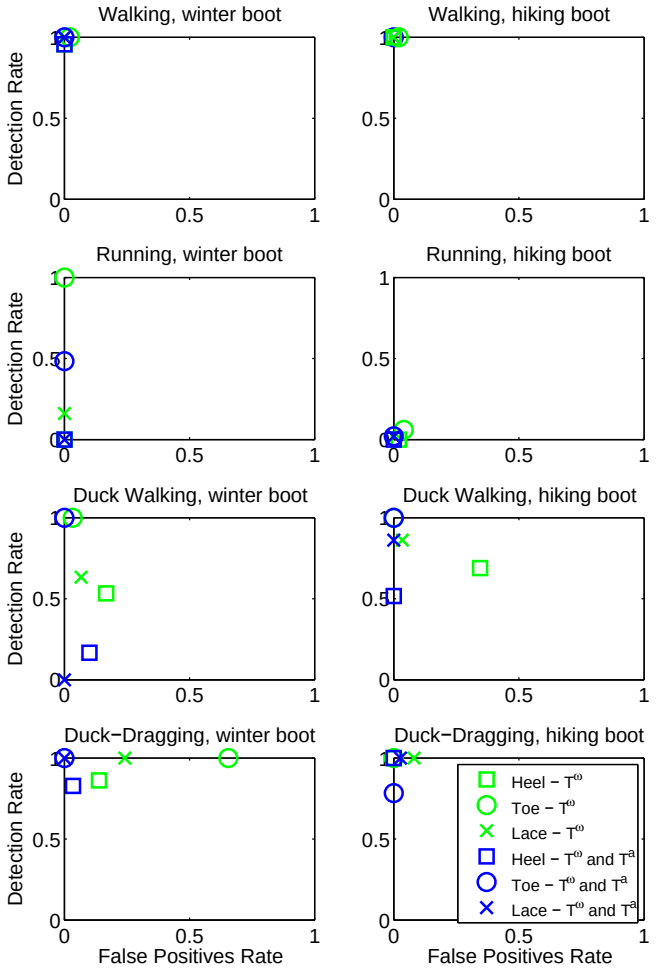

Figure 4.8: *Stand still detection results using the winter boot (left column) and the hiking boot (right column) for the 99% probability threshold. Top row is walking experiments, second is running, third is duck walking, last is duck-dragging. Heel mounted sensor is marked with squares, toe mounted with circles and crosses marks lace mounted sensor. The best corner to be in is the top left combining a good stand still detection ability with few false positives. Overall, the toe mounted sensor using the combination of angular velocity and acceleration measurements seems to provide the most reliable results.*

detected, especially for a toe mounted sensor in which case almost all would then have been detected. When looking at the data the reason for this is clear. While running, angular velocity wise the boot is almost never really still. During each stance the rigid boot is rolling against the ground why the sensor is always moving. A stand still is therefore impossible to detect since it in reality does not happen. This is probably true also for higher sampling frequencies. Further studies are needed to establish this.

The heel and lace mounted sensors have significant problems detecting stand stills even when the softer winter boot is used. Not even the accelerometer provides reliable results indicating that those parts of the boot are rarely stationary while running.

Duck Walking

The third movement, duck walking, is walking in a squatted position. The heels never touch the ground making the results open for discussion. A heel mounted sensor should therefore never be stationary during a duck walking sequence. Sometimes though, the foot is still enough to indicate a stand still even though the heel is not touching the ground. We have though chosen to list these as true positives. The same reasoning is also valid for the lace mounted sensor.

The stand still performance during duck walking is shown in the third row of Figure 4.8 and in Table 4.4. The toe mounted sensor is the most reliable which is natural since the toes are the only part that actually touch the ground. The rigidness of the hiking boot also makes it easier to detect stand stills using the lace mounted sensor compared to using the winter boot.

Duck-drag

Duck-drag is a movement where the whole foot with the sensor is flat on the ground with the other knee on the ground. The movement is common among firefighters while searching a room with poor visibility.

The last row of in Figure 4.8 and Table 4.5 shows the stand still detection results for a duck-drag sequence. Detecting the stand stills is straight forward no matter the sensor position or the boot. The whole foot is flat on the ground during each stance why they are readily detected. The problems are the false positives. Using only the gyro results in many false positives if the softer boot is used. Supporting the gyro measurements with the accelerometer signals significantly reduces the false positives in this case.

Using a rigid boot, the heel mounted sensor performs the best. The toe mounted sensor misses some true positives but it has no false positives. The lace mounted sensor causes a couple of false positives.

4.2.2 Discussion

The problem of reliable stand still detection using a foot mounted IMU is not a straightforward one. We have in this section shown that both the type of boot and the position of the sensor will greatly affect the stance detection performance.

The worst sensor position seems to be a heel mounted sensor. During many movements the heel hardly ever touches the ground making the stand stills nearly impossible to detect. In some cases the stand stills are impossible to detect since they do not even exist because the heel is always moving. Simply falling for the temptation to hide the sensor in the heel just because it appears easy, seems to be a poor choice.

The lace mounted sensor performs quite well during most movements. The most problematic movement is running since the laces are very rarely still while running. It is better to put the sensor by the laces than by the heel, but it does not seem to be the ultimate sensor position.

Putting the sensor by the toes appears to be the best position. Stand stills have been shown to be safely detected using almost all movements and boot combinations. It seems though that the gyro based stand still detection needs some support by either the accelerometer or some logical rules to reduce the number of false positives. A very rigid boot can still make the stand stills difficult to detect during running which must be noted if a positioning system is to be built for police officers or soldiers.

Walking sequence			
Winter boot, 46 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	44	46	46
False positives	0	1	0
Stand stills detected, T^{a} and T^{ω}	44	46	46
False positives	0	0	0
Hiking boot, 50 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	50	50	50
False positives	0	1	0
Stand stills detected, T^{a} and T^{ω}	50	50	50
False positives	0	0	0

Table 4.2: *Experimental results of the stand still detection performance of a walking sequence using different sensor positions on two different boots.*

Running sequence			
Winter boot, 31 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	0	31	5
False positives	0	0	0
Stand stills detected, T^{a} and T^{ω}	0	15	0
False positives	0	0	0
Hiking boot, 40 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	0	3	1
False positives	1	2	0
Stand stills detected, T^{a} and T^{ω}	0	1	1
False positives	0	0	0

Table 4.3: *Stand still performance during a running sequence using different sensor positions.*

Duck Walking sequence			
Winter boot, 30 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	16	30	19
False positives	5	1	2
Stand stills detected, T^{a} and T^{ω}	5	30	0
False positives	3	0	0
Hiking boot, 29 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	20	29	25
False positives	10	0	1
Stand stills detected, T^{a} and T^{ω}	15	29	25
False positives	0	0	0

Table 4.4: *Stand still detection performance during a duck walking sequence. The toe mounted sensor is the most reliable.*

4.3 Localization Experiments

Localization experiments where the position of a walking subject was estimated have been performed. The system uses the foot mounted IMU and zero velocity updates based on the stand still detection framework described in Section 4.1.

4.3.1 Measurements

Three types of measurements are used in the localization experiments. First, IMU data is used to estimate the movements of the boot. Second, virtual measurements of zero velocity are created when a stand still event has been detected. Third, the experiments were performed in the LiU VICON lab which provides accurate position measurements that are used as ground truth to evaluate the system performance.

Duck-drag sequence			
Winter boot, 29 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	25	29	29
False positives	4	19	7
Stand stills detected, T^{a} and T^{ω}	24	29	29
False positives	1	0	0
Hiking boot, 37 steps	Heel	Toe	Lace
Stand stills detected, T^{ω}	37	37	37
False positives	0	0	3
Stand stills detected, T^{a} and T^{ω}	37	29	37
False positives	0	0	1

Table 4.5: *Stand still detection performance during a duck-drag sequence. The softer winter boot gives plenty of false positives when only gyro data is used to detect the stand stills.*

IMU data

The IMU sensor provides acceleration and gyro measurements which are used to calculate how the boot is moved and to detect stand still phases. The acceleration measurements $\mathbf{y}_{a,t}$ are three dimensional and contain both user movements and earth gravity. To remove the earth gravity component from the accelerometer data, the orientation of the sensor must be known. A small error in orientation will result in gravity components being misinterpreted as user accelerations.

Angular velocity is measured using the three dimensional gyro measurements $\mathbf{y}_{\omega,t}$. They are used to keep track of the orientation of the boot. By simply adding the gyro measurements, the orientation can be estimated. Unfortunately, a long term heading drift will be introduced since only relative heading measurements are used.

The sensor also provides measurements of the magnetic field. If only the earth magnetic field was measured it could be used as a compass measurement. Unfortunately, indoors the magnetic field is disturbed by steel structures and electrical wiring making the measurements unreliable. These disturbances are commonly stronger close to the floor which is bad news for a foot mounted magnetometer. We have started to study methods to verify that the magnetometer is reliable based on gyro studies, but it has not been used in these experiments. Consequently, no magnetometer is used in the localization experiments.

Virtual Stand Still Measurements

Each time a stand still has been detected, a virtual measurement of zero velocity is used to update the filter estimates. A small noise is added to represent the uncertainty in the stand still detection.

Position Ground Truth

The undertaken positioning experiments were performed in the VICON lab at LiU to obtain a ground truth to compare the estimated position with. The size of the active area in the lab is about 3 by 5 meters making the experiments quite limited in size. Therefore the user has to walk back and forth in the area to produce meaningful experiments.

4.3.2 Models

To estimate the position of a user based on accelerations, knowledge of the orientation of the sensor is crucial to compute the user movements. The orientation is represented using quaternions, $\mathbf{q}_t$, which are four dimensional. These are updated by the gyro measurements, used as direct inputs in the time update.

The position must also be represented by states since that is what we truly want to estimate. Position is tracked in three dimensions giving the three states $\mathbf{p}_t$. Since the filtering framework uses virtual zero velocity measurements, also velocity states are needed so the measurements can be incorporated correctly, giving three more states $\mathbf{v}_t$. Both velocity and position are updated using acceleration measurements used as direct inputs in the time update.

The ten dimensional state vector is therefore

$$\mathbf{x}_t = \begin{pmatrix} \mathbf{p}_t & \mathbf{v}_t & \mathbf{q}_t \end{pmatrix}^T . \tag{4.16}$$

Dynamical Model

The dynamical model used is as mentioned above on input form where the accelerometer and gyro measurements are used as direct inputs in the model.

$$\begin{aligned} \begin{pmatrix} \mathbf{p}_{t+T} \\ \mathbf{v}_{t+T} \\ \mathbf{q}_{t+T} \end{pmatrix} &= \begin{pmatrix} I & TI & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix} \begin{pmatrix} \mathbf{p}_t \\ \mathbf{v}_t \\ \mathbf{q}_t \end{pmatrix} + \begin{pmatrix} \frac{T^2}{2} I \\ TI \\ 0 \end{pmatrix} \Big(R^T(\mathbf{q}_t)\big(\mathbf{y}_{a,t} + \mathbf{r}_{a,t}\big) + \mathbf{w}_{a,t} \Big) \\ &+ \frac{T}{2} \begin{pmatrix} 0 \\ 0 \\ S'(\mathbf{q}_t) \end{pmatrix} \big(\mathbf{y}_{\omega,t} + \mathbf{r}_{\omega,t} + \mathbf{w}_{\omega,t} \big) \end{aligned} \tag{4.17}$$

where I is a 3×3 identity matrix. $R^T(\mathbf{q}_t)$ rotates the local measurements of acceleration to global acceleration measurements using the quaternion estimates. $S'(\mathbf{q}_t)$ converts the gyro measurements into changes in orientation. These conversions are described more thoroughly in Appendix A.

There are four parameters representing the uncertainties of the model. $\mathbf{r}_a$ and $\mathbf{r}_\omega$ are measurement noise from the accelerometer and gyro, respectively. Added to these are $\mathbf{w}_a$ and $\mathbf{w}_\omega$ which are process noises added to the position and velocity covariances and orientation covariances, respectively. These measurement and process noise terms do not represent the same uncertainties. In the acceleration case, the measurement noise represents the uncertainties in the acceleration measurement used as input while the process noise term represents the model errors

introduced by assuming that the acceleration is constant over the time interval. The two noise terms can of course be added to create a third new noise term replacing the two old ones. Correspondingly, $\mathbf{w}_\omega$ represents the errors introduced by assuming constant angular velocity. It must therefore be converted from three dimensional uncertainties in orientation to four dimensional quaternion covariances.

Measurement Model

Since no measurements of position, velocity or orientation exist, the only measurements available are the virtual measurements of zero velocity whenever a stand still has been detected.

$$\mathbf{0} = \mathbf{v}_t + \mathbf{r}_{0,t}. \tag{4.18}$$

A slight noise $\mathbf{r}_{0,t}$ has been added to represent the uncertainty in the stand still detection framework.

4.3.3 Filter

Since a nonlinearity is present due to the rotation matrices, an extended Kalman filter is used to estimate the states, see Algorithm 2 Section 2.3.2. It mainly consists of a time update step since the measurement update only enters when a stand still has been detected.

Time Update

The time update of the state estimates uses $\mathbf{y}_a$ and $\mathbf{y}_\omega$ according to

$$\begin{pmatrix} \hat{\mathbf{p}}_{t+T|t} \\ \hat{\mathbf{v}}_{t+T|t} \\ \hat{\mathbf{q}}_{t+T|t} \end{pmatrix} = \begin{pmatrix} I & TI & 0 \\ 0 & I & 0 \\ 0 & 0 & I \end{pmatrix} \begin{pmatrix} \hat{\mathbf{p}}_{t|t} \\ \hat{\mathbf{v}}_{t|t} \\ \hat{\mathbf{q}}_{t|t} \end{pmatrix} + \begin{pmatrix} \frac{T^2}{2} I \\ TI \\ 0 \end{pmatrix} R^T(\hat{\mathbf{q}}_{t|t})\mathbf{y}_{a,t} + \frac{T}{2} \begin{pmatrix} 0 \\ 0 \\ S'(\hat{\mathbf{q}}_{t|t}) \end{pmatrix} \mathbf{y}_{\omega,t} \tag{4.19}$$

to update the states. The quaternion estimates must thereafter be normalized

$$\hat{\mathbf{q}}_{t+T|t+T} := \frac{\hat{\mathbf{q}}_{t+T|t+T}}{\|\hat{\mathbf{q}}_{t+T|t+T}\|} \tag{4.20}$$

to ensure that they are still unit quaternions. The time update of the state covariances P depend on the jacobians of the dynamical model f where f_x' are the jacobians of (4.17) with respect to the states and f_v' are the jacobians of (4.17) with respect to the noise.

$$P_{t+T|t} = f_x'(\hat{\mathbf{x}})P_{t|t}f_x'(\hat{\mathbf{x}})^T + f_v'(\hat{\mathbf{x}})Q_v f_v'(\hat{\mathbf{x}})^T \tag{4.21}$$

where the state jacobians are

$$f_x'(\hat{\mathbf{x}}) = \begin{pmatrix} I & TI & \frac{T^2}{2}\frac{\partial R^T(\hat{\mathbf{q}}_{t|t})}{\partial \hat{\mathbf{q}}_{t|t}}\mathbf{y}_{a,t} \\ 0 & I & T\frac{\partial R^T(\hat{\mathbf{q}}_{t|t})}{\partial \hat{\mathbf{q}}_{t|t}}\mathbf{y}_{a,t} \\ 0 & 0 & I + \frac{T}{2}\frac{\partial S'(\hat{\mathbf{q}}_{t|t})}{\partial \hat{\mathbf{q}}_{t|t}}\mathbf{y}_{\omega,t} \end{pmatrix} \tag{4.22}$$

and the noise jacobians are

$$f'_v(\hat{\mathbf{x}}) = \begin{pmatrix} \frac{T^2}{2}R^T(\hat{\mathbf{q}}_{t|t}) & \frac{T^2}{2}I & 0 & 0 \\ TR^T(\hat{\mathbf{q}}_{t|t}) & TI & 0 & 0 \\ 0 & 0 & \frac{T}{2}S'(\hat{\mathbf{q}}_{t|t}) & \frac{T}{2}S'(\hat{\mathbf{q}}_{t|t}) \end{pmatrix}. \tag{4.23}$$

The process noise matrix is

$$Q_v = \begin{pmatrix} \sigma_a^2 I & 0 & 0 & 0 \\ 0 & \sigma_{w_a}^2 I & 0 & 0 \\ 0 & 0 & \sigma_\omega^2 I & 0 \\ 0 & 0 & 0 & \sigma_{w_\omega}^2 I \end{pmatrix} \tag{4.24}$$

where σ_a^2, $\sigma_{w_a}^2$, σ_ω^2 and $\sigma_{w_\omega}^2$ are the noise covariances of the acceleration measurement noise, the acceleration process noise, the gyro measurement noise and the angular velocity process noise, respectively.

Measurement Update

The measurement update occurs only when a stand still has been detected. For the state and covariance updates, the jacobians of the measurement function with respect to the states, h'_x, and noise, h'_{r_0}, are needed.

$$h'_x = \begin{pmatrix} 0 & I & 0 \end{pmatrix} \tag{4.25}$$

$$h'_{r_0} = I \tag{4.26}$$

The Kalman gain K can now be calculated as

$$\begin{aligned} S &= h'_x P_{t+T|t} h'^T_x + h'_{r_0} R h'^T_{r_0} \\ K &= P_{t+T|t} h'^T_x S^{-1}. \end{aligned} \tag{4.27}$$

where $R = r_0 I$. The state estimates and covariances are now updated according to

$$\begin{aligned} \hat{\mathbf{x}}_{t+T|t+T} &= \hat{\mathbf{x}}_{t+T|t} - K\left(\mathbf{0} - \hat{\mathbf{v}}_{t+T|t}\right) \\ P_{t+T|t+T} &= P_{t+T|t} - K h'_x P_{t+T|t}. \end{aligned} \tag{4.28}$$

After the measurement update the quaternion estimates must once again be normalized.

4.3.4 Experimental Results

A couple of walking experiments have been performed for which ground truth is available. Since the experiments had to be performed on quite a limited area due to the VICON system, the area was traversed several time to make the experiments large enough. The downside is that when position estimate and ground truth are plotted together, the plots become quite messy.

Experiment 1

The first experiment covers three laps of the area totalling roughly 36 meters during 35 seconds, Figure 4.9. The accumulated error in heading is surprisingly small and the uncertainty estimates covers the ground truth within two standard deviations. The x-, y- and z-components with two standard deviations are also shown individually in Figure 4.10.

Experiment 2

The second experiment is slightly longer, about 42 meters, and covers multiple sharp corners, Figure 4.11. These 180° turns are underestimated resulting in a heading misalignment causing the estimates to drift off. The estimated uncertainties are though accurate and keeps the ground truth within two standard deviations. Again, the estimated x-, y- and z-positions are plotted with ground truth and two standard deviations in Figure 4.12.

4.3.5 Discussion and Future Work

The positioning framework described above suffers from two shortcomings. First of all there are no heading measurements, making it impossible to recover from accumulated heading errors. The gyros are not perfect why these errors are prone to happen. One solution could be to incorporate a magnetometer in the measurements and thereby receive heading estimates. The problem is as described before that the magnetometer measurements are unreliable due to plenty of corrupting disturbances. We have started to look at using the gyro measurements to decide when the magnetometer is reliable and when it is not. By detecting that the magnetometer measurements are reliable and only incorporating these into the estimate, the heading errors could be removed. As of today, the reliance test works quite well in the one dimensional case.

The second problem is the lack of position measurements or relations between position estimates. If some sort of loop closure measurements were available the performance of the positioning would be greatly improved. We will in the near future look into using for example magnetometer disturbances or some sort of active markers as a mean to create loop closure measurements.

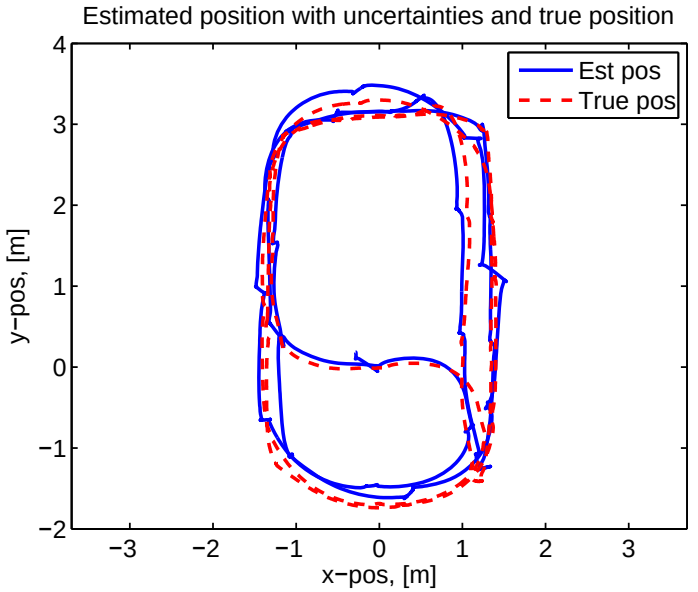

Figure 4.9: *First positioning experiment with estimate, solid line, and ground truth, dashed line. The experiment started and ended in the origin.*

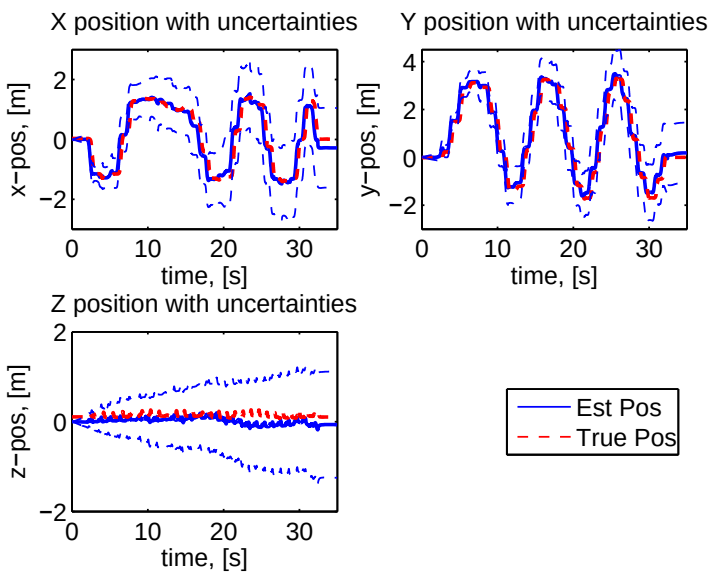

Figure 4.10: *x-, y- and z-position estimates and ground truth of the first positioning estimate. 2σ of estimate uncertainties are shown in thin dashed line.*

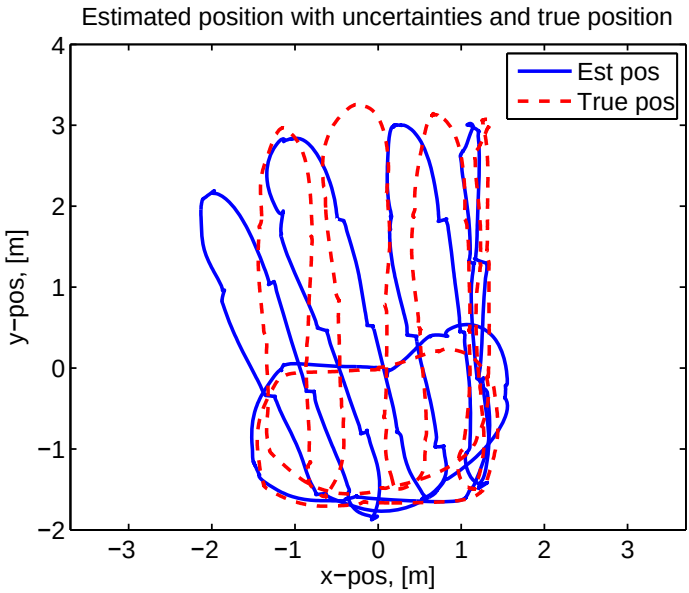

Figure 4.11: *Second positioning experiment with estimate, solid line, and ground truth, dashed line. The experiment started and ended in the origin.*

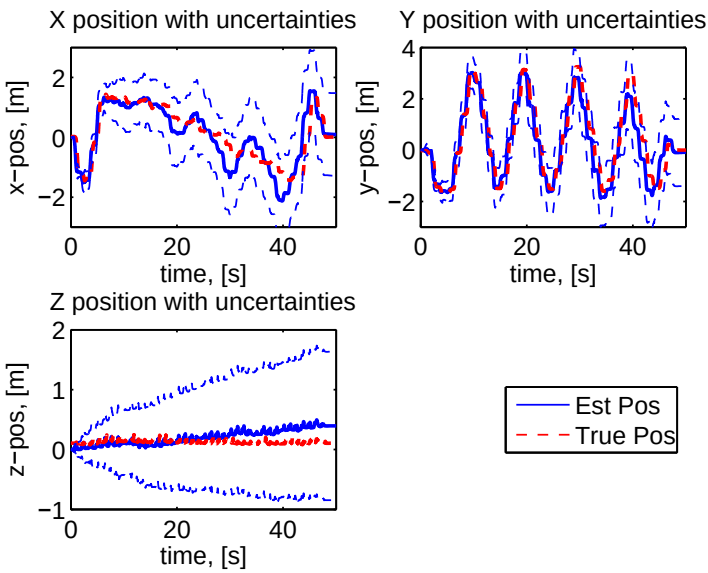

Figure 4.12: *x-, y- and z-position estimates with ground truth for experiment 2. 2σ of estimate uncertainties are shown in thin dashed line.*

5

Conclusions and Future Work

5.1 Conclusions

The conclusions have previously been drawn under their respective sections but some will be repeated here for completeness.

5.1.1 Indoor Localization

The work on indoor localization was divided into two parts. A stand still detection framework was presented which was evaluated using different boots and sensor positions. The stand still detection framework was incorporated in a localization experiment with very accurate ground truth.

The stand still detection framework is based on either gyro signals only or a combination of gyro and accelerometer signals. The stand still distributions of the signals are used in conjunction with a Hidden Markov Model to estimate the probabilities of stand still and movement. Using both the measured accelerations and angular velocities to detect that the foot is at rest has been shown to provide the safest stand still detections. Still some false positives do occur and a basic requirement of forcing the boot to be still for at least k measurements, $k \sim \{2, 3\}$, would probably enhance the performance.

The IMU should preferably be mounted near the toes on the boot. This area is almost always stationary during stand still no matter the movement why it seems to be a better sensor position than the heel. The more rigid the boot, the more difficult it is to detect the stand stills, most notably because the sensor is never really at rest since the boot rolls over the ground due to its sturdiness.

The localization framework suffers from the absence of heading and position mea-

surements. The ZUPT experiments provide fairly good results when compared to ground truth but large turns often tend to be underestimated causing the position estimate to drift off.

5.1.2 RADAR SLAM

A navigation framework for surface vessels using RADAR measurements only was presented in Paper A. The system was intended to be used during GPS failure and is shown to provide accurate navigation results during an extensive experiment. The navigation framework treats the RADAR scans like images and uses the SIFT algorithm to track features between scans. These provide information about how the vessel is moving and turning which was used to estimate the vessel position.

5.1.3 Underwater Sensor Positioning

A passive underwater sensor localization scheme is presented in Paper B where triaxial magnetometers and a friendly vessel with known magnetic characteristics is used to determine the sensors positions. Simulations indicate that if the vessel is equipped with a GNSS a positioning error of 12.9% can be achieved. The simulations also indicate that our positioning scheme is quite insensitive to minor errors in sensor orientation and magnetic signature, when GNSS is used throughout the trajectory.

5.2 Future Work

The future work will primarily be undertaken in indoor localization. That said, the other work presented in this thesis could be quite useful in this task as well. If the outcome of the future work described below is that one or two IMUs are not enough to do accurate firefighter positioning, other sensors will be needed too. Range measurements from a SONAR or possibly even a RADAR sensor might be needed to assist the IMU positioning framework. Another possibility is to map the magnetic characteristics of the environment in a SLAM fashion and incorporate this into the positioning using loop closures.

The first steps will though just involve IMUs. Can we expand the magnetometer reliance test mentioned earlier to also work in three dimensions? Incorporating this would greatly enhance the localization performance since the heading error is the most troubling one.

The problem of tracking crawling firefighters has not been solved though just because walking ones can be tracked. A second IMU probably needs to be put on the knee to detect when it touches the ground. The foot mounted one will most likely not be still enough while crawling, why a knee mounted sensor is probably needed. This will be looked into soon.

Finally, as mentioned earlier, incorporating some sort of loop closure measurements, preferably ones based on the way firefighters search a building, would allow great positioning improvements. The system should be designed using the

fundamental difference between the human sensor platform and the robotic platform; a thinking person. By exploiting the human in the system and designing the positioning solution appropriately and carefully, the human sensor platform can be just as informative as the robotic one, if not more. We have ideas for how this can be done and they will be implemented in the near future.

A

Quaternion Properties

Quaternions were discovered by Hamilton as an extension for the imaginary numbers into three dimensions, Hamilton (1844). Later, the unit quaternion started to be used for angle representation in rotations providing singularity free rotations. For a thorough description of quaternions see Kuipers (1999).

A.1 Operations and Properties

A quaternion is a four-tuple of real numbers denoted by $q = (q_0, q_1, q_2, q_3)$. Alternatively it can be described as consisting of a scalar part q_0 and the vector $\mathbf{q}$.

$$q = \begin{pmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{pmatrix} = \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix} \tag{A.1}$$

Quaternion multiplication is denoted as $\odot$ and defined as

$$p \odot q = \begin{pmatrix} p_0 \\ \mathbf{p} \end{pmatrix} \odot \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix} = \begin{pmatrix} p_0 q_0 - \mathbf{p} \cdot \mathbf{q} \\ p_0 \mathbf{q} + q_0 \mathbf{p} + \mathbf{p} \times \mathbf{q} \end{pmatrix}. \tag{A.2}$$

Some quaternions properties are:

$$p \odot q \neq q \odot p \tag{A.3}$$

$$\text{norm}(q) = \sqrt{\sum_{i=0}^{3} q_i^2} \tag{A.4}$$

$$q^{-1} = \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix}^{-1} = \begin{pmatrix} q_0 \\ -\mathbf{q} \end{pmatrix} \tag{A.5}$$

The unit quaternion used for rotation operations also fulfills $\text{norm}(q) = 1$.

A.2 Describing a Rotation using Quaternions

The quaternion

$$q = \begin{pmatrix} \cos\delta \\ \sin\delta\mathbf{n} \end{pmatrix} \tag{A.6}$$

describes a rotation around the vector $\mathbf{n}$ with the angle 2δ. There are two ways of depicting a rotation; either the coordinate frame is rotated and the vector is fixed or the vector is rotated and the coordinate frame is fixed. The difference is in the sign of the rotation. In this work the vector will be assumed constant and the coordinate system is rotated.

Given the vectors

$$v = \begin{pmatrix} 0 \\ \mathbf{v} \end{pmatrix} \quad \text{and} \quad u = \begin{pmatrix} 0 \\ \mathbf{u} \end{pmatrix} \tag{A.7}$$

a rotation of v around $\mathbf{n}$ can be written as

$$u = q^{-1} \odot v \odot q \tag{A.8}$$

which is the assumed standard rotation in this work. The resulting rotation is

$$v = \begin{pmatrix} \mathbf{q} \cdot \mathbf{u} q_0 - (q_0\mathbf{u} - \mathbf{q} \times \mathbf{u}) \cdot \mathbf{q} \\ (\mathbf{q} \cdot \mathbf{u})\mathbf{q} + q_0(q_0\mathbf{u} - \mathbf{q} \times \mathbf{u}) + (q_0\mathbf{u} - \mathbf{q} \times \mathbf{u}) \times \mathbf{q} \end{pmatrix} \tag{A.9}$$

which simplifies to

$$v = \begin{pmatrix} 0 \\ 2(\mathbf{q} \cdot \mathbf{u})\mathbf{q} + (q_0^2 - \mathbf{q} \cdot \mathbf{q})\mathbf{u} - 2q_0\mathbf{q} \times \mathbf{u} \end{pmatrix}. \tag{A.10}$$

A.3 Rotation Matrix

The quaternion rotation (A.10) can be rewritten as a matrix multiplication

$$\mathbf{v} = R(q)\mathbf{u} \tag{A.11}$$

where

$$R(q) = \begin{pmatrix} (q_0^2 + q_1^2 - q_2^2 - q_3^2) & 2(q_1 q_2 + q_0 q_3) & 2(q_1 q_3 - q_0 q_2) \\ 2(q_1 q_2 - q_0 q_3) & (q_0^2 - q_1^2 + q_2^2 - q_3^2) & 2(q_2 q_3 + q_0 q_1) \\ 2(q_1 q_3 + q_0 q_2) & 2(q_2 q_3 - q_0 q_1) & (q_0^2 - q_1^2 - q_2^2 + q_3^2) \end{pmatrix}. \tag{A.12}$$

A.4 Quaternion Dynamics

In the case of the quaternions describing a rotation between a global coordinate system and a local one attached to a moving sensor, the description of the quaternions will contain some dynamics. The full derivation of the quaternion dynamics can be studied in Törnqvist (2008) but parts will be recited here.

Let the quaternion q_{lg} represent the rotation of the local coordinate system in the global one. The angular velocity of the sensor unit in the local coordinate system is ω_{lg}^l which can be written as

$$\omega_{lg}^l = \begin{pmatrix} 0 \\ \omega_x \\ \omega_y \\ \omega_z \end{pmatrix} = \begin{pmatrix} 0 \\ \boldsymbol{\omega} \end{pmatrix} \tag{A.13}$$

while q_{lg} is

$$q_{lg} = \begin{pmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{pmatrix} = \begin{pmatrix} q_0 \\ \mathbf{q} \end{pmatrix}. \tag{A.14}$$

The quaternion derivative can now be written as

$$\begin{aligned} \dot{q}_{lg} &= \frac{1}{2} q_{lg} \odot \omega_{lg}^l = \frac{1}{2} \begin{pmatrix} -\mathbf{q} \cdot \boldsymbol{\omega} \\ q_0 \boldsymbol{\omega} + \mathbf{q} \times \boldsymbol{\omega} \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} -(q_1 \omega_x + q_2 \omega_y + q_3 \omega_z) \\ q_0 \boldsymbol{\omega} - \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix} \begin{pmatrix} q_1 \\ q_2 \\ q_3 \end{pmatrix} \end{pmatrix} \\ &= \underbrace{\frac{1}{2} \begin{pmatrix} 0 & -\omega_x & -\omega_y & -\omega_z \\ \omega_x & 0 & \omega_z & -\omega_y \\ \omega_y & -\omega_z & 0 & \omega_x \\ \omega_z & \omega_y & -\omega_x & 0 \end{pmatrix}}_{S(\omega_{lg}^l)} \begin{pmatrix} q_0 \\ q_1 \\ q_2 \\ q_3 \end{pmatrix} = \underbrace{\frac{1}{2} \begin{pmatrix} -q_1 & -q_2 & -q_3 \\ q_0 & -q_3 & q_2 \\ q_3 & q_0 & -q_1 \\ -q_2 & q_1 & q_0 \end{pmatrix}}_{S'(q_{lg})} \begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix} \end{aligned}$$

If the angular velocity ω_t is assumed constant over the sampling interval the noise free discrete time model is

$$q_{t+1} = e^{\frac{1}{2}S(\omega_t)T} q_t \tag{A.15}$$

where a Taylor series expansion gives

$$\begin{aligned} e^{\frac{1}{2}S(\omega_t)T} &= \sum_{n=0}^{\infty} \frac{\left(\frac{1}{2}S(\omega_t)T\right)^n}{n!} \\ &= \sum_{n=0}^{\infty} \left(\frac{T}{2}\right)^n \frac{1}{n!} S(\omega_t)^n. \end{aligned} \tag{A.16}$$

If the sampling time T is short, the expansion can be approximated with the first two terms

$$e^{\frac{1}{2}S(\omega_t)T} \approx I + \frac{T}{2} S(\omega_t) \tag{A.17}$$

giving the discrete time model

$$\begin{aligned} q_{t+1} &= q_t + \frac{T}{2} S(\omega_t) q_t \\ &= q_t + \frac{T}{2} S'(q_t)\omega_t. \end{aligned} \tag{A.18}$$

If angular velocity noise term ν_ω is included the discrete model becomes

$$q_{t+1} = q_t + \frac{T}{2} S'(q_t)\omega_t + \frac{T}{2} S'(q_t)\nu_{\omega,t}. \tag{A.19}$$

Bibliography

P. Aggarwal, D. Thomas, L. Ojeda, and J. Borenstein. Map matching and heuristic elimination of gyro drift for personal navigation systems in gps-denied conditions. *Measurement Science and Technology*, 22(2):025205, 2011.

S. Beauregard. Omnidirectional pedestrian navigation for first responders. In *Proc. of the 4th Workshop on Positioning, Navigation and Communication, WPNC07*, Hannover, Germany, 2007.

J. Callmer, K. Granström, J. Nieto, and F. Ramos. Tree of words for visual loop closure detection in urban slam. In *Proceedings of the 2008 Australasian Conference on Robotics and Automation (ACRA)*, 2008.

J. Callmer, M. Skoglund, and F. Gustafsson. Silent localization of underwater sensors using magnetometers. *EURASIP Journal on Advances in Signal Processing*, 2010a.

J. Callmer, D. Törnqvist, and F. Gustafsson. Probabilistic stand still detection using foot mounted IMU. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010b.

J. Callmer, D. Törnqvist, H. Svensson, P. Carlbom, and F. Gustafsson. Radar SLAM using visual features. *EURASIP Journal on Advances in Signal Processing*, 2010c. Under revision.

The Economist. No jam tomorrow. *The Economist Technology Quarterly*, 398 (8724):20–21, 2011.

R. Feliz, E. Zalama, and J. G. Garcia-Bermejo. Pedestrian tracking using inertial sensors. *Journal of Physical Agents*, 3(1):35–43, 2009.

E. Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.

S. Godha, G. Lachapelle, and M. E. Cannon. Integrated GPS/INS system for pedestrian navigation in signal degraded environment. In *Proc. of ION GNSS*, 2006.

N.J. Gordon, D.J. Salmond, and A.F.M. Smith. A novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings on Radar and Signal Processing*, 140:107–113, 1993.

K. Granström, J. Callmer, F. Ramos, and J. Nieto. Learning to detect loop closure from range data. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2009.

A. Grant, P. Williams, N. Ward, and S. Basker. GPS jamming and the impact on maritime navigation. *Journal of Navigation*, 62(2):173–187, 2009.

S. Grzonka, F. Dijoux, A. Karwath, and W. Burgard. Mapping indoor environments based on human activity. In *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, Anchorage, Alaska, 2010.

F. Gustafsson. *Statistical Sensor Fusion*. Studentlitteratur, 2010.

S.W.R. Hamilton. On quaternions; or on a ner system of imaginaries in algebra. *Philosophical Magazine*, xxv:10–13, 1844.

A.R. Jiménez, F. Seco, J.C. Prieto, and J. Guevara. Indoor pedestrian navigation using an ins/ekf framework for yaw drift reduction and a foot-mounted imu. In *Positioning Navigation and Communication (WPNC), 2010 7th Workshop on*, pages 135 –143, march 2010.

S.J. Julier, J.K. Uhlmann, and H.F. Durrant-Whyte. A new approach for filtering nonlinear systems. In *American Control Conference, 1995. Proceedings of the*, volume 3, pages 1628 –1632, June 1995.

R.E. Kalman. A new approach to linear filtering and prediction problems. *Trans. ASME J Basic Engr.*, 82:35–45, 1960.

J.B. Kuipers. *Quaternions and Rotation Sequences*. Princeton University Press, 1999.

F. Lindsten, J. Callmer, H. Ohlsson, D. Törnqvist, T. B. Schön, and F. Gustafsson. Geo-referencing for UAV navigation using environmental classification. In *Proceedings of 2010 International Conference on Robotics and Automation (ICRA)*, 2010.

L. Ojeda and J. Borenstein. Non-GPS navigation for security personnel and first responders. *Journal of Navigation*, 60(3):391–407, 2007.

J. Rantakokko, P. Strömbäck, J. Rydell, J. Callmer, D. Törnqvist, F. Gustafsson, P. Händel, M. Jobs, and M. Grudén. Accurate and reliable soldier and first responder indoor positioning: Multi-sensor systems and cooperative localization. *IEEE Wireless Communications Magazine*, 2011. Accepted for publication.

P. Robertson, M. Angermann, and B. Krach. Simultaneous localization and mapping for pedestrians using only foot-mounted inertial sensors. In *Proceedings of the 11th international conference on Ubiquitous computing*, Ubicomp '09, pages 93–96, 2009.

W. J. Rugh. *Linear System Theory*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition, 1996.

I. Skog. *Low-cost Navigation Systems*. PhD thesis, KTH, Stockholm, Sweden, 2009.

D. Törnqvist. *Estimation and Detection with Applications to Navigation*. PhD thesis, Linköping University, Linköping, Sweden, 2008.

N. Wahlström, J. Callmer, and F. Gustafsson. Magnetometers for tracking metallic targets. In *Proceedings of the International Conference on Information Fusion (FUSION)*, 2010.

N. Wahlström, J. Callmer, and F. Gustafsson. Single target tracking using vector magnetometers. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011.

Widyawan, M. Klepal, and S. Beauregard. A backtracking particle filter for fusing building plans with pdr displacement estimates. In *Positioning, Navigation and Communication, 2008. WPNC 2008. 5th Workshop on*, pages 207 –212, march 2008.

O. Woodman and R. Harle. Rf-based initialisation for inertial pedestrian tracking. In *Proceedings of the 7th International Conference on Pervasive Computing*, 2009.

Licentiate Theses
Division of Automatic Control
Linköping University

P. Andersson: Adaptive Forgetting through Multiple Models and Adaptive Control of Car Dynamics. Thesis No. 15, 1983.

B. Wahlberg: On Model Simplification in System Identification. Thesis No. 47, 1985.

A. Isaksson: Identification of Time Varying Systems and Applications of System Identification to Signal Processing. Thesis No. 75, 1986.

G. Malmberg: A Study of Adaptive Control Missiles. Thesis No. 76, 1986.

S. Gunnarsson: On the Mean Square Error of Transfer Function Estimates with Applications to Control. Thesis No. 90, 1986.

M. Viberg: On the Adaptive Array Problem. Thesis No. 117, 1987.

K. Ståhl: On the Frequency Domain Analysis of Nonlinear Systems. Thesis No. 137, 1988.

A. Skeppstedt: Construction of Composite Models from Large Data-Sets. Thesis No. 149, 1988.

P. A. J. Nagy: MaMiS: A Programming Environment for Numeric/Symbolic Data Processing. Thesis No. 153, 1988.

K. Forsman: Applications of Constructive Algebra to Control Problems. Thesis No. 231, 1990.

I. Klein: Planning for a Class of Sequential Control Problems. Thesis No. 234, 1990.

F. Gustafsson: Optimal Segmentation of Linear Regression Parameters. Thesis No. 246, 1990.

H. Hjalmarsson: On Estimation of Model Quality in System Identification. Thesis No. 251, 1990.

S. Andersson: Sensor Array Processing; Application to Mobile Communication Systems and Dimension Reduction. Thesis No. 255, 1990.

K. Wang Chen: Observability and Invertibility of Nonlinear Systems: A Differential Algebraic Approach. Thesis No. 282, 1991.

J. Sjöberg: Regularization Issues in Neural Network Models of Dynamical Systems. Thesis No. 366, 1993.

P. Pucar: Segmentation of Laser Range Radar Images Using Hidden Markov Field Models. Thesis No. 403, 1993.

H. Fortell: Volterra and Algebraic Approaches to the Zero Dynamics. Thesis No. 438, 1994.

T. McKelvey: On State-Space Models in System Identification. Thesis No. 447, 1994.

T. Andersson: Concepts and Algorithms for Non-Linear System Identifiability. Thesis No. 448, 1994.

P. Lindskog: Algorithms and Tools for System Identification Using Prior Knowledge. Thesis No. 456, 1994.

J. Plantin: Algebraic Methods for Verification and Control of Discrete Event Dynamic Systems. Thesis No. 501, 1995.

J. Gunnarsson: On Modeling of Discrete Event Dynamic Systems, Using Symbolic Algebraic Methods. Thesis No. 502, 1995.

A. Ericsson: Fast Power Control to Counteract Rayleigh Fading in Cellular Radio Systems. Thesis No. 527, 1995.

M. Jirstrand: Algebraic Methods for Modeling and Design in Control. Thesis No. 540, 1996.

K. Edström: Simulation of Mode Switching Systems Using Switched Bond Graphs. Thesis No. 586, 1996.

J. Palmqvist: On Integrity Monitoring of Integrated Navigation Systems. Thesis No. 600, 1997.
A. Stenman: Just-in-Time Models with Applications to Dynamical Systems. Thesis No. 601, 1997.
M. Andersson: Experimental Design and Updating of Finite Element Models. Thesis No. 611, 1997.
U. Forssell: Properties and Usage of Closed-Loop Identification Methods. Thesis No. 641, 1997.
M. Larsson: On Modeling and Diagnosis of Discrete Event Dynamic systems. Thesis No. 648, 1997.
N. Bergman: Bayesian Inference in Terrain Navigation. Thesis No. 649, 1997.
V. Einarsson: On Verification of Switched Systems Using Abstractions. Thesis No. 705, 1998.
J. Blom, F. Gunnarsson: Power Control in Cellular Radio Systems. Thesis No. 706, 1998.
P. Spångéus: Hybrid Control using LP and LMI methods – Some Applications. Thesis No. 724, 1998.
M. Norrlöf: On Analysis and Implementation of Iterative Learning Control. Thesis No. 727, 1998.
A. Hagenblad: Aspects of the Identification of Wiener Models. Thesis No. 793, 1999.
F. Tjärnström: Quality Estimation of Approximate Models. Thesis No. 810, 2000.
C. Carlsson: Vehicle Size and Orientation Estimation Using Geometric Fitting. Thesis No. 840, 2000.
J. Löfberg: Linear Model Predictive Control: Stability and Robustness. Thesis No. 866, 2001.
O. Härkegård: Flight Control Design Using Backstepping. Thesis No. 875, 2001.
J. Elbornsson: Equalization of Distortion in A/D Converters. Thesis No. 883, 2001.
J. Roll: Robust Verification and Identification of Piecewise Affine Systems. Thesis No. 899, 2001.
I. Lind: Regressor Selection in System Identification using ANOVA. Thesis No. 921, 2001.
R. Karlsson: Simulation Based Methods for Target Tracking. Thesis No. 930, 2002.
P.-J. Nordlund: Sequential Monte Carlo Filters and Integrated Navigation. Thesis No. 945, 2002.
M. Östring: Identification, Diagnosis, and Control of a Flexible Robot Arm. Thesis No. 948, 2002.
C. Olsson: Active Engine Vibration Isolation using Feedback Control. Thesis No. 968, 2002.
J. Jansson: Tracking and Decision Making for Automotive Collision Avoidance. Thesis No. 965, 2002.
N. Persson: Event Based Sampling with Application to Spectral Estimation. Thesis No. 981, 2002.
D. Lindgren: Subspace Selection Techniques for Classification Problems. Thesis No. 995, 2002.
E. Geijer Lundin: Uplink Load in CDMA Cellular Systems. Thesis No. 1045, 2003.
M. Enqvist: Some Results on Linear Models of Nonlinear Systems. Thesis No. 1046, 2003.
T. Schön: On Computational Methods for Nonlinear Estimation. Thesis No. 1047, 2003.
F. Gunnarsson: On Modeling and Control of Network Queue Dynamics. Thesis No. 1048, 2003.
S. Björklund: A Survey and Comparison of Time-Delay Estimation Methods in Linear Systems. Thesis No. 1061, 2003.

M. Gerdin: Parameter Estimation in Linear Descriptor Systems. Thesis No. 1085, 2004.

A. Eidehall: An Automotive Lane Guidance System. Thesis No. 1122, 2004.

E. Wernholt: On Multivariable and Nonlinear Identification of Industrial Robots. Thesis No. 1131, 2004.

J. Gillberg: Methods for Frequency Domain Estimation of Continuous-Time Models. Thesis No. 1133, 2004.

G. Hendeby: Fundamental Estimation and Detection Limits in Linear Non-Gaussian Systems. Thesis No. 1199, 2005.

D. Axehill: Applications of Integer Quadratic Programming in Control and Communication. Thesis No. 1218, 2005.

J. Sjöberg: Some Results On Optimal Control for Nonlinear Descriptor Systems. Thesis No. 1227, 2006.

D. Törnqvist: Statistical Fault Detection with Applications to IMU Disturbances. Thesis No. 1258, 2006.

H. Tidefelt: Structural algorithms and perturbations in differential-algebraic equations. Thesis No. 1318, 2007.

S. Moberg: On Modeling and Control of Flexible Manipulators. Thesis No. 1336, 2007.

J. Wallén: On Kinematic Modelling and Iterative Learning Control of Industrial Robots. Thesis No. 1343, 2008.

J. Harju Johansson: A Structure Utilizing Inexact Primal-Dual Interior-Point Method for Analysis of Linear Differential Inclusions. Thesis No. 1367, 2008.

J. D. Hol: Pose Estimation and Calibration Algorithms for Vision and Inertial Sensors. Thesis No. 1370, 2008.

H. Ohlsson: Regression on Manifolds with Implications for System Identification. Thesis No. 1382, 2008.

D. Ankelhed: On low order controller synthesis using rational constraints. Thesis No. 1398, 2009.

P. Skoglar: Planning Methods for Aerial Exploration and Ground Target Tracking. Thesis No. 1420, 2009.

C. Lundquist: Automotive Sensor Fusion for Situation Awareness. Thesis No. 1422, 2009.

C. Lyzell: Initialization Methods for System Identification. Thesis No. 1426, 2009.

R. Falkeborn: Structure exploitation in semidefinite programming for control. Thesis No. 1430, 2010.

D. Petersson: Nonlinear Optimization Approaches to $\mathcal{H}_2$-Norm Based LPV Modelling and Control. Thesis No. 1453, 2010.

Z. Sjanic: Navigation and SAR Auto-focusing in a Sensor Fusion Framework. Thesis No. 1464, 2011.

K. Granström: Loop detection and extended target tracking using laser data. Thesis No. 1465, 2011.