From Posture to Motion: The Challenge for Real Time Wireless Inertial Motion Capture.

A. D. Young
ayoung9@inf.ed.ac.uk

Institute for Computing Systems Architecture
School of Informatics, University of Edinburgh
10 Crichton Street, Edinburgh EH8 9AB, United Kingdom

ABSTRACT

Networks of Wireless Inertial Measurement Units (WIMUs) allow for the real time capture of orientation data from multiple devices. Combining data from WIMUs with a rigid body model allows estimation of subject posture. However, posture information is not sufficient for full-body motion capture, the position of the subject in space must also be tracked. In this paper we examine the problem of relative position tracking; provide a demonstration of a novel location estimation algorithm, based on probabilistic ground contact detection; and discuss the use of aiding technologies.

Categories and Subject Descriptors

C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems

General Terms

Theory, Algorithms

Keywords

Wireless body sensor network, inertial motion capture, real-time, Kalman filtering

1. INTRODUCTION

Wireless inertial posture tracking [15] allows for the real time reconstruction of subject posture from a body-worn network of Wireless Inertial Measurement Units (WIMUs). The orientations of WIMUs are estimated relative to a global Earth-fixed co-ordinate frame. By mapping the orientations of WIMUs to bones in a rigid body model the posture of the subject can be reconstructed, as illustrated in Figure 1. This reconstruction preserves spatial relationships between bones in the body, but does not provide information about the location of the body in space.

BodyNets'10, September 10–12, 2010, Corfu, Greece.

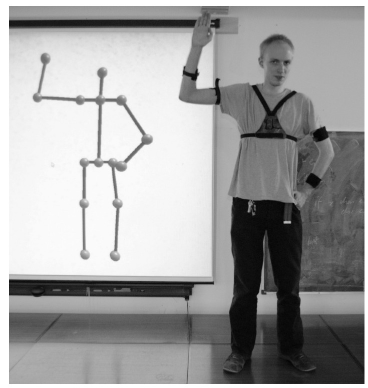

Figure 1: Example of real time posture reconstruction from a network of WIMUs.

While posture information alone allows for the development of many applications [2,1,9], full motion capture is desirable to compete with alternative motion capture technologies, such as acoustic, magnetic and optical systems. Unlike existing technologies, inertial systems present the possibility of full-body motion capture without the requirement for external infrastructure or line-of-sight links [13]. Freedom from infrastructure requirements allows wireless inertial motion capture to be performed in almost any location, potentially reducing costs, as no dedicated capture space is required, and lowering the barrier to entry in areas such as animation, healthcare, and human computer interaction.

This paper investigates the problem of moving from posture only capture to full motion capture, including estimation of relative location. The problem of directly deriving position information from WIMUs is discussed and shown to be inherently flawed. An alternative method of position tracking, based on full-body posture and ground contact information is proposed and demonstrated to allow for full-body motion capture.

The remainder of the paper is structured as follows: Section 2 positions the paper in relation to existing commercial and academic works; Section 3 discusses the constraints imposed by body sensor network hardware limitations; Section 4 presents the theory of spatial tracking using WIMUs; Section 5 provides preliminary simulations of full-body spatial tracking; and finally, Section 6 presents the conclusions of the work.

2. RELATED WORK

We consider two areas of development in inertial motion capture: commercial wired motion capture suits, and academic research into GPS-free location systems.

Three commercial inertial motion capture systems have been developed: *IGS-190-M* by Animazoo, *3DSuit* by Innalabs, Inc. and *MVN* by Xsens Technologies B.V. All of these systems use a wired network of inertial/magnetic sensors mounted in a body suit. Data from the network is gathered to a central unit for wireless transmission to a PC for processing. A comparison of the three systems is presented in [7], concluding that the *MVN* system has the highest accuracy and ability. Details of the *MVN* system are presented in an Xsens white paper [12]. The system uses a probabilistic kinematic body model, floor contact information, and Kalman filtering to estimate subject motion from redundant sensor data.

In academic research the focus has typically been on tracking a subject's movement through space rather than full motion capture. Numerous approaches involving infrastructure placed throughout a building have been developed using technologies such as Ultra-Wide Band, Received Signal Strength Indicator, and Ultrasound. The use of infrastructure is useful in that it provides a reference frame for drift-free tracking but is also a disadvantage as deployed infrastructure limits the available capture volume.

Work on inertial dead-reckoning [17, 10, 14] has typically focussed on the use of a single foot-worn IMU to track successive strides. Such systems use double integration of IMU accelerometer data, combined with zero-velocity update techniques, to estimate position. As with the method of position tracking to be proposed in this paper, these solutions, relying purely on dead-reckoning, suffer from ever increasing position estimation drift. The advantage of the integration approaches is that they implicitly handle motions where ground contact is lost, such as running. The disadvantage is that, only having a single foot sensor, existing integration approaches are unable to handle more complex motions such as rolling.

3. LIMITATIONS OF BODY SENSOR NETWORK HARDWARE

In order to build small, light-weight, unobtrusive devices, Body Sensor Network (BSN) nodes are constrained to use low-power processors and radios. This in turn limits the available computational power and network bandwidth.

In contrast to existing commercial systems, the *Orient* inertial capture system, developed in previous work [15], is fully wireless. In order to achieve tracking of multiple devices with a single low-bandwidth radio channel the orientation of *Orient* devices is estimated locally on each device and the resulting orientation quaternion transmitted to a central PC for body model processing. Operating in this manner allows a reduction in data rate of 79% compared to transmitting unprocessed sensor samples. Our current wireless implementation, using a typical 250 kbit/s low-power radio, is capable of tracking the orientations of fifteen WIMUs, enough to reconstruct the major segments of the human body, at 64 Hz each using a Time Division Multiple Access (TDMA) network.

The data transfer requirement for the *Orient* system compared to a typical WIMU, the InterSense InertiaCube-3 [6], that transfers raw sensor data, is shown in Table 1. The reduction in data rate, produced by pre-processing data locally on each WIMU, is the key to supporting the simultaneous multiple device capture required for motion capture.

Table 1: Data transmission requirements for raw data transfer and *Orient* system

Method	Data transferred	Data rate per device
Raw data	9×12bit@180 Hz	19440 bit/s
Orient	4×16bit@64 Hz	4096 bit/s

4. TRACKING THEORY

A typical WIMU consists of three tri-axial sensors: accelerometers, to measure linear accelerations, including gravity; rate gyroscopes, to measure angular velocity; and magnetometers, to measure the Earth's magnetic field.

The orientation of the sensor is estimated by integrating the outputs of the rate gyroscopes and combining this with an estimate derived from observing the position of the accelerometer and magnetometer vectors [3]. By combining data from two different orientation estimates in this manner a drift-free estimate of the three Degree of Freedom (3DoF) orientation can be formed.

In order to provide 6DoF tracking, including an estimate of the translation of the WIMU, it is necessary to perform double integration of acceleration data. A linear accelerometer measures both the static acceleration due to gravity and the dynamic acceleration due to motion. It is therefore necessary to remove the gravitational acceleration from the measured acceleration vector prior to integration. The ability to cancel the effects of gravity depends on the ability of the WIMU to estimate its static and dynamic orientation with a high degree of accuracy. An error as low as 1 mrad results in an acceleration of approximately 0.01 m/s^2 which, left uncorrected, would result in an error of 4.5 m after 30 s [13]. Such levels of error are considerably lower than the quoted static accuracy of low cost IMUs and will certainly be exceeded due the effect of linear accelerations on the orientation estimation process during motion [16]. Without a mechanism to correct for this drift it is clear that accurate 6DoF tracking of WIMUs is a hopeless endeavour.

4.1 Zero Velocity Update

In order to counteract the unbounded increase in error due to accelerometer error integration the technique of Zero Velocity Update (ZVU) is commonly applied. ZVU operates by zeroing the velocity estimate when the IMU in known to be stationary. In simple applications zeroing the velocity helps reduce the build up of error while the device is stationary. In more complex applications, employing Kalman filters, ZVU allows the filter to estimate bias conditions [5, Ch. 13], further increasing estimation accuracy.

In order to apply ZVU it is necessary to know when an IMU is expected to be stationary. For the tracking of human locomotion it is known that, except during a slip, the contact point between the foot and the ground is stationary during the stance phase of the gait cycle. This assumption is valid for all gaits including walking, running and jumping. Position estimation is therefore performed by double integration of acceleration data over each stride, with ZVU applied while the foot is stationary.

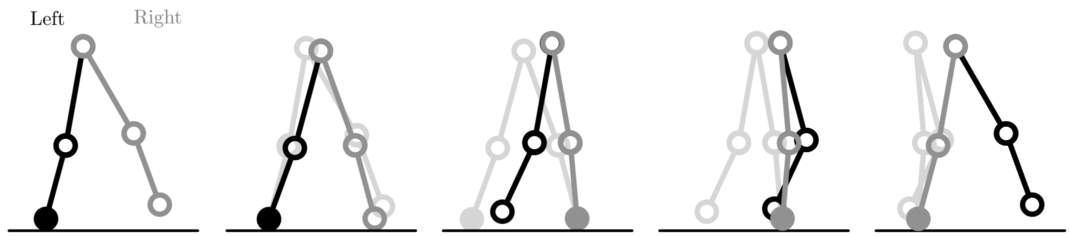

(a) Left foot anchor Right leg terminal swing
(b) Anchor transition Double stance
(c) Right foot anchor Left leg initial swing
(d) Right foot anchor Left leg mid swing
(e) Right foot anchor Left leg terminal swing

Figure 2: Example of lowest-point algorithm applied to walking gait of a simplified body model.

While the contact point between the foot and the floor is stationary during the stance phase, the foot itself is not stationary due to its rolling motion [11, Ch. 2]. ZVU is therefore performed as the leg is in the mid-stance phase [10] as the subject's body weight moves over the stationary foot. This phase occurs in the interval of approximately 10 to 30% of the gait cycle, measured from the initial contact point of the foot touching the ground. Identification of the onset of the stance phase is generally performed by detecting the sharp acceleration caused by the heel striking the floor [17]. The onset of mid-stance is then generally estimated from knowledge of the nominal gait cycle period.

The requirement to estimate the onset of the mid-stance gait phase, based on detection of a heel strike, limits the applicability of the ZVU approach to subjects with normal gait function. For subjects whose gait does not include a pronounced heel strike, such as those suffering from motor-control illnesses or deformities, the ZVU approach is liable to fail as the zero velocity state cannot be reliably recognised. This difficulty extends to other unusual gait patterns, such as tip-toeing, that may be desirable in animation applications.

4.2 Model-Based Contact Tracking

We now consider the problem of tracking position based primarily on the orientation of the subject's limb segments. Applying the orientations of WIMUs to associated joints in the rigid body model of the subject allows the posture of the subject to be estimated. Our method of tracking position is to attempt to track the contact between the subject and the ground. We choose to investigate position tracking using only orientation data as sending additional data would require an increase in data transfer requirements for each device, resulting in a corresponding decrease in network update rate.

We initially assume that one joint in the rigid body model, called the *anchor*, is attached to the ground at each frame. The positions of all other joints in the model can then be determined in relation to the fixed position of the anchor. In order to achieve motion it is necessary that at some point in time the selected anchor joint is changed. We shall now consider various approaches for correctly selecting the anchor joint.

4.2.1 Lowest Point Algorithm

The simplest algorithm for selecting the anchor joint is to assume that the ground level is at a constant height and that the joint with the lowest height is attached to the ground.

```
# Find current lowest point in the model
lowest = anchor
for joint in model.joints:
    if joint.height < lowest.height:
        lowest = joint

if not lowest == anchor:
    # Update absolute position of anchor
    # from difference in relative positions
    # between old and new anchors
    aPos += lowest.pos - anchor.pos
    # Lock anchor position to ground level
    aPos.height = 0
    # Update anchor reference
    anchor = lowest

# Update absolute position of root joint
# relative to absolute anchor position
rootPosition = aPos - anchor.pos
```

Listing 1: Lowest point algorithm

The lowest point algorithm, presented as Listing 1, is run on every frame after the calculation of the position of each joint relative to the root of the model. The algorithm first finds the lowest joint in the model. If this joint is lower than the current anchor then the lowest joint becomes the new anchor. The absolute position of the anchor is fixed at the instant that the anchor transition occurs. The absolute position of the root of the body model is calculated based on the fixed absolute position of the anchor and the position of the anchor relative to the root.

The operation of the lowest point algorithm applied to walking gait is illustrated in Figure 2. To begin with the left foot is anchored to the ground with the right leg approaching terminal swing. The transition occurs just after the right leg makes contact with the ground and the left leg is about start its swing. After the transition the right foot is anchored and the hips move over the anchor to continue their forward motion.

Despite its simplicity, the lowest point algorithm is capable of producing lifelike animations provided the constraint of a constant ground level is observed. As the ground contact point is not constrained to being the feet, the algorithm is capable of capturing a surprising variety of motions including rolling and hand-stands. However, the assumption of a level ground surface prevents the lowest point algorithm from tracking more complex motions such as walking on uneven surfaces or climbing steps. Additionally, the requirement for continual contact with the ground precludes the algorithm from tracking gaits such as jogging and running.

Determining the contact point between the body model and the ground based on the lowest point heuristic is prone to error when two joints are in close proximity to the ground. For example, during the mid swing gait phase, Figure 2d, the swinging foot passes close to the ground and minor errors in either WIMU orientation estimation or body model proportions may lead to the swinging foot appearing to be momentarily lower than the planted foot. The lowest point algorithm is unable to filter out this transitory condition, resulting in a rapid, and visually unappealing, change in hip velocity.

4.2.2 Contact Aiding Sensors

We now briefly consider the use of aiding sensors to improve detection of ground contact.

An obvious solution for normal gait capture would be to add pressure sensors to the subject's shoes, for example a pressure sensitive insole. The use of a pressure sensitive pad, and suitable thresholding algorithm, would allow a simple binary indication of ground contact to be transmitted by each foot WIMU. As this would require only a single extra bit of information to be transmitted it would be possible to include within the existing TDMA networking schedule.

The addition of accurate contact information would allow the restriction of a constant ground level to be relaxed. Removing the constant ground level assumption would allow tracking to be performed on uneven surfaces including stairs. However, the requirement of continuous contact with the ground would remain.

While the use of aiding sensors appears initially appealing, it has two obvious problems: cost and generality. The increase in cost to provide contact sensing, while minor in comparison to total system cost, should not be overlooked. The requirement to provide an electrical connection between a pressure-sensitive insole, or other aiding sensor, to each foot WIMU presents an additional failure point.

The greater problem with the use of aiding sensors is how to deal with non-gait motions. As discussed in Section 4.2.1, the lowest point algorithm makes no assumptions as to which joint is in contact with the ground, resulting in the ability to track a greater variety of motion. Contact sensors are necessarily constrained to pre-determined contact points. It is unclear how this limitation could be overcome.

4.2.3 Hidden Markov Models

An alternative to attempting direct measurement of contact is to use a Hidden Markov Model (HMM) to describe the gait cycle, where states include information pertaining to ground contact. This approach has been demonstrated successfully for estimating the point of contact between the foot and the ground for subjects with normal gait function using an ear-worn accelerometer [8].

If normal bipedal gait was the only form of locomotion, HMMs would be a good solution as considerable information is available from the subject posture that could be used to infer hidden model state. However, the human body is capable of many other forms of locomotion such as rolling and crawling. These alternatives, which differ considerably in their patterns of ground contact, would require their own model structures. The resulting composite models are liable to become extremely complex. While the structure of such models could be learned, provided sufficient annotated data were available, the HMM approach would be liable to failure when faced with new situations.

4.2.4 Summary of Contact Tracking

Ground contact tracking provides a simple mechanism for estimating subject motion under the assumption that at least one point of the body is in contact with the ground at any instant. However, the approaches for detecting ground contact presented so far are limited in their ability to perform under all circumstances. The lowest point algorithm, while capable of tracking many motions, is sensitive to errors in orientation estimation and assumes a constant ground level. The use of contact sensors or HMMs requires assumptions to be made as to the possible contact points, limiting their generality. The common assumption of constant ground contact precludes tracking of many interesting motions such as running and jumping.

4.3 Contact Tracking Kalman Filter

Two major problems were identified with the basic contact tracking algorithms: firstly, the difficulty in inferring ground contact points in a general manner; secondly, the limitation to constant ground contact.

We first investigate a method for inferring ground contact in a reliable and generalised manner. To do this we employ Newton's first law of motion that states that:

> Every body persists in its state of being at rest or of moving uniformly straight forward, except insofar as it is compelled to change its state by force impressed [4].

Recall that the simple lowest point algorithm was flawed in that errors in orientation estimation of subject body model could result in sudden changes in hip velocity. These sudden changes, unless accompanied by sufficient force, violate Newton's law. From Newton's second law we know that a force results in an acceleration that can be measured by an accelerometer. We therefore propose to augment the sensed data with the acceleration of the hips. This requires only a single extra data vector to be transmitted with minimal effect to network performance. Alternatively, if processing of model translation were to be performed in situ on the body, the hip IMU could be made into a special, more powerful, device, thereby removing its contribution to local network data rate.

We can now formulate a general rule for selecting the most likely ground contact point. Reformulating the position update steps shown in Listing 1 it can be shown that the velocity of the anchor joint in the hip-fixed co-ordinate frame, derived from subsequent frames of posture data, is exactly the opposite of the hip velocity in the world co-ordinate frame. Using this knowledge we can calculate the set of all possible velocities, $\{\boldsymbol{v}_j : j \in 0, 1, \ldots, N\}$, corresponding to the hypothesis that joint j is stationary at the current instant.

Selection of the appropriate anchor joint can then be achieved by:

$$j = \operatorname*{argmin}_{j \in N} \left(\|\hat{\boldsymbol{v}} - \boldsymbol{v}_j\| \right), \tag{1}$$

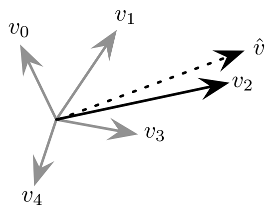

Figure 3: Selection of possible velocities based on predicted velocity.

where $\hat{\boldsymbol{v}}$ is the best estimate of the hip velocity for the current frame. This process is illustrated in Figure 3, where $\boldsymbol{v}_2$ would be selected as the new velocity measurement.

4.3.1 Kalman Filter Design

The method of selecting the anchor joint based on predicted subject velocity requires that velocity at each frame can be estimated accurately. This may be achieved through the use of a Kalman filter. For the purposes of initial experiments we select the state vector to comprise the velocity and acceleration of the subject's hip:

$$\boldsymbol{x} = [v_x, v_y, v_z, a_x, a_y, a_z]^T . \tag{2}$$

The state update equation, in partitioned matrix form, is defined to be:

$$\boldsymbol{x}_k = \boldsymbol{A}\boldsymbol{x}_{k-1} + \boldsymbol{w}_k = \begin{bmatrix} \boldsymbol{I} & \boldsymbol{\Delta T} \\ \boldsymbol{0} & \boldsymbol{C} \end{bmatrix} \boldsymbol{x}_{k-1} + \boldsymbol{w}_k, \tag{3}$$

where $\boldsymbol{I}$ is the 3×3 identity matrix, $\boldsymbol{0}$ is the 3×3 zero matrix, $\boldsymbol{\Delta T}$ is the 3×3 diagonal matrix $\mathrm{diag}\left([dt, dt, dt]\right)$ where dt is the time between updates, $\boldsymbol{C}$ is the 3×3 diagonal matrix $\mathrm{diag}\left(\left[e^{-\frac{dt}{\tau}}, e^{-\frac{dt}{\tau}}, e^{-\frac{dt}{\tau}}\right]\right)$ where τ is a time constant resulting in an exponential decay, and $\boldsymbol{w}$ is a process noise vector with covariance matrix $\boldsymbol{Q}$. The state update equation predicts the hip acceleration as a coloured noise signal and the velocity to be determined by integration of acceleration data.

As the state variables can be directly observed, the relationship between state variables and observed measurements, $\boldsymbol{z}$, is defined as:

$$\boldsymbol{z}_k = \boldsymbol{H}\boldsymbol{x}_k + \boldsymbol{v}_k = \boldsymbol{x}_k + \boldsymbol{v}_k, \tag{4}$$

where $\boldsymbol{v}$ is a measurement noise vector with covariance matrix $\boldsymbol{R}$.

The hip acceleration can be observed directly from the hip IMU. However, the velocity of the hip must be estimated based on the velocity of an anchor joint. Selection of an appropriate anchor joint is performed using Equation 1 where the a priori velocity estimate of the Kalman filter is used as $\hat{\boldsymbol{v}}$.

The selection of the current hip velocity measurement based on predicted velocity presents a problem when there are several possible candidates. Selection of the wrong estimate may result in an increase in estimate inaccuracy, that may ultimately result in filter divergence. This issue requires considerable further investigation to establish its severity. One area of investigation would be the use of a multiple hypothesis filter [5, Ch. 3].

The position of the subject, relative to the initial position, is estimated by integration of the Kalman filtered velocity estimate.

4.3.2 Detecting Loss of Ground Contact

All the ground contact tracking solutions discussed so far have been limited to tracking motions where at least one point of the body model remains in constant contact with the ground at any instant. This limitation prevents such approaches from tracking motions such as running and jumping.

The Kalman filter structure allows for variable weighting of observed measurements through the measurement noise covariance matrix $\boldsymbol{R}$. If the hypothesised velocity measurement can be rejected then the velocity measurement can be ignored. We select as the null hypothesis:

H_{0_j}: Joint j was stationary at iteration k.

The statistic we choose to use for the test is the probability that the velocity $\boldsymbol{v}_j$ was drawn from the probability distribution specified by the a priori expected velocity, $\hat{\boldsymbol{v}}^-$, and covariance, $\boldsymbol{\Sigma}$, estimated by the Kalman filter prediction. The deviation from the predicted value is calculated as the Mahalnobis distance:

$$d = \sqrt{\left(\boldsymbol{v}_j - \hat{\boldsymbol{v}}^-\right)^T \boldsymbol{\Sigma}^{-1} \left(\boldsymbol{v}_j - \hat{\boldsymbol{v}}^-\right)}. \tag{5}$$

The two-tailed p-value for this distance is then calculated and the probability of H_{0_j}, i.e. the probability of $\boldsymbol{v}_j$ being observed, estimated as:

$$P\left(H_{0_j}\right) = P(\boldsymbol{v}_j|\hat{\boldsymbol{v}}^-) = 2 \cdot \Phi\left(-|d|\right), \tag{6}$$

where Φ is the standard normal cumulative distribution function.

The significance level for rejecting the null hypothesis is selected to balance the risk of Type I error — rejecting the velocity update when it is in fact valid — and Type II error — using the velocity update when it is erroneous. In the case of the Kalman filter presented here, Type II errors are likely to be more serious, as these provide misleading information, and so the probability of rejecting the null hypothesis should be set fairly high.

The probability of the null hypothesis can also be used as an alternative method for the anchor selection. We can replace Equation 1 with:

$$j = \underset{j \in N}{\mathrm{argmax}} \left(P(H_{0_j})\right). \tag{7}$$

This results in selecting as the anchor the joint which we have least reason to suspect was in motion.

5. DEMONSTRATION

In order to demonstrate operation of the proposed tracking algorithms they were simulated using data from the Carnegie Mellon University motion capture database. The optical motion capture data was processed to produce simulated inertial sensor data, including linear accelerations, angular velocities and magnetic field vector data. Use of optical motion capture data provided a truth reference against which the various tracking algorithms could be compared. For the purposes of preliminary demonstration the algorithms were simulated without the addition of sensor measurement noise or orientation estimation inaccuracy.

Four algorithms were simulated:

Int double integration of hip IMU estimated acceleration,

LP the lowest point algorithm shown in Listing 1,

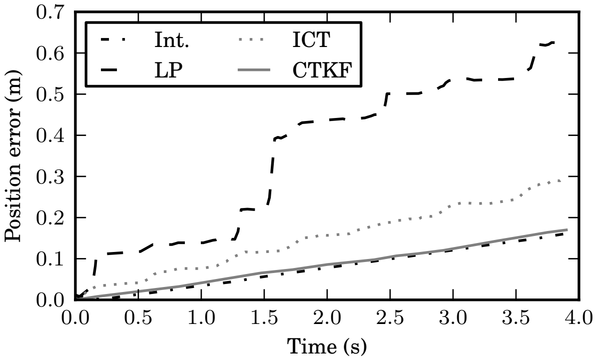

(a) Walk (CMU Subject 16 trial 15.)

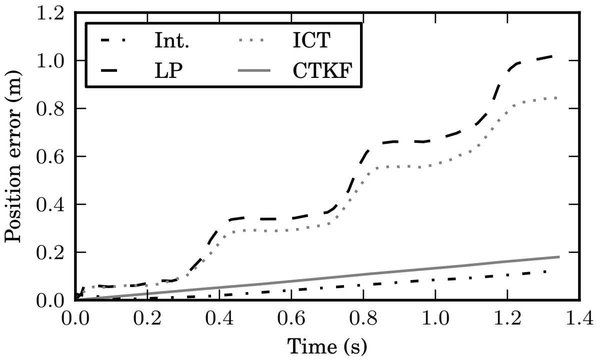

(b) Run (CMU Subject 16 trial 35.)

Figure 4: Position error between optical reference and estimated position.

ICT idealised contact tracker where ground contact is determined using Equation 1, with $\hat{\boldsymbol{v}}$ taken as being the true hip velocity,

CTKF Kalman filter structure described in Section 4.3.1 with anchor selection performed according to Equation 7.

The ability of the algorithms to estimate the position of a walking subject is illustrated in Figure 4a. The error is calculated as the magnitude of the difference between the estimated position and true position taken from the optical motion capture.

The lowest point algorithm is seen to perform the worst with sudden jumps in error caused by mistaken selection of anchor joint. The idealised anchor selection algorithm also displays small jumps in error. These errors are ascribed to the fact that the model joints do not directly correspond to the true stationary points. This results in systematic errors in the resulting velocity estimates.

The Kalman filter and integration approaches show very similar behaviour. This is to be expected in the idealised simulation without significant measurement noise corrupting sensed data. Simulations including sensor noise indicate that Kalman filter will require further tuning to accurately track subject motion. This process is left as the subject for further work.

The ability of the contact tracking Kalman filter to identify stationary joints is demonstrated in Figure 5a where the pattern of left and right foot ground contact, including the brief double stance period, can be clearly observed. Figure 5b illustrates the ground contact probabilities for a running subject. The characteristic of running, the gap between footfalls where the body loses contact with the ground, can be seen clearly.

The utility of the content tracking Kalman filter is illustrated in Figure 4b. The lowest point and idealised contact tracker, which both assume constant contact, suffer from substantial errors on each stride. The Kalman filter, by ignoring erroneous ground contact, is able to integrate acceleration data over the flight phase of the running gait.

6. CONCLUSIONS

Moving from posture estimation to full motion tracking of a subject in space represents a significant challenge for wireless inertial sensing. The limitations of low power body sensor networks, with limited radio bandwidth, require the use of novel algorithms that limit the amount of information that must be transmitted within the network.

In order to reduce data transfer requirements the problem of tracking relative subject motion based primarily on subject posture has been addressed. The use of ground contact information has been identified as the key to translating postural data to locomotion.

A novel algorithm, combining a Kalman filter with a probabilistic method of inferring ground contact has been developed. This algorithm provides an advantage over ground contact sensors, or Markov models, as it makes no assumptions about possible contact locations. This makes the new algorithm more flexible in the range of motions it can capture. The ability of the new algorithm to track subject motion and infer ground contact has been demonstrated for walking and running gaits.

While the proposed algorithm shows significant potential, further work is required to tune Kalman filter operation to optimally weight available sensor data. The tuning of Kalman state update and covariance matrices, along with investigation of multiple-hypothesis filtering is the goal of ongoing research.

Acknowledgements

This work was supported by the UK Engineering and Physical Sciences Research Council under the Basic Technology Research Programme, Grant C523881. The optical motion data used in this project was obtained from `http://mocap.cs.cmu.edu`. The database was created with funding from NSF EIA-0196217.

7. REFERENCES

[1] Arvind, D., and Bartosik, M. Clustering of motion data from on-body wireless sensor networks for human-imitative walking in bipedal robots. In *Advanced Robotics, Proceedings of the 14th International Conference on* (June 2009), pp. 1–6.

[2] Arvind, D., and Bates, A. The speckled golfer. In *Body Area Networks, Proceedings of the 3rd international conference on* (2008), p. 28.

[3] Bachmann, E. R. *Inertial and Magnetic Tracking of Limb Segment Orientation for Inserting Humans into Synthetic Environments*. PhD thesis, Naval Postgraduate School, Monterey, California, 2000.

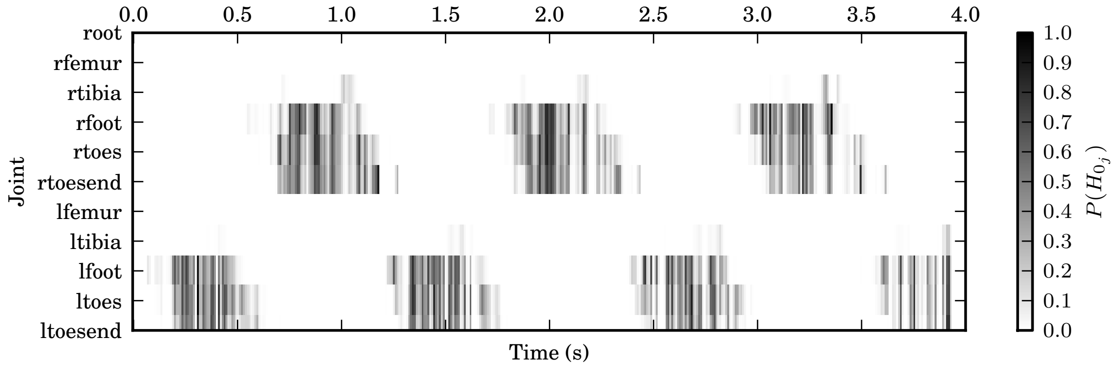

(a) Walk (CMU Subject 16 trial 15)

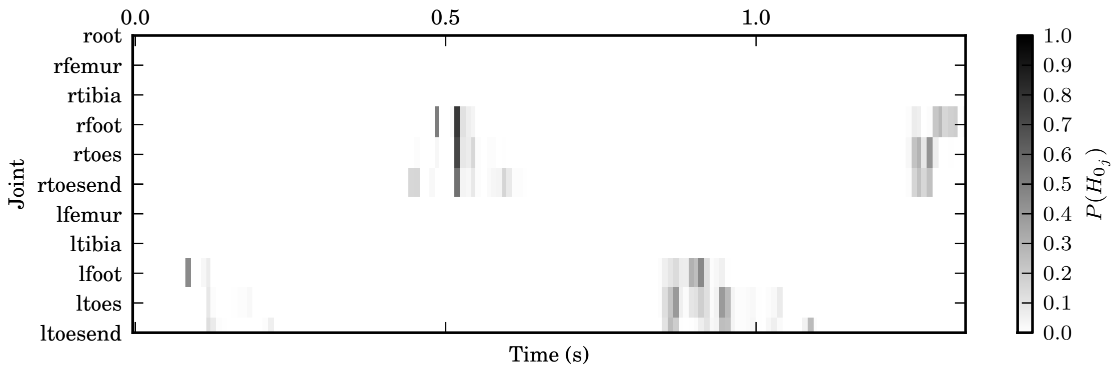

(b) Run (CMU Subject 16 trial 35)

Figure 5: Ground contact hypothesis probabilities for walking and running gaits.

[4] Cohen, I., and Whitman, A. *The Principia, A new translation.* University of California press, 1999.

[5] Groves, P. D. *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems.* Artech House, 2008.

[6] InterSense, Inc. *InertiaCube3 Datasheet.*

[7] Kider Jr, J., Stocker, C., and Badler, N. Untethered motion capture evaluation for flightline maintenance support. Tech. rep., Department of Computer and Information Science, University of Pennsylvania, June 2008.

[8] Lo, B., Yang, G.-Z., and Pansiot, J. Bayesian analysis of sub-plantar ground reaction force with BSN. In *Wearable and Implantable Body Sensor Networks, International Workshop on* (June 2009), IEEE Computer Society, pp. 133 –137.

[9] Lympouridis, V., Parker, M., Young, A., and Arvind, D. Sonification of Gestures Using Specknets. In *Sound and Music Computing Conference, Proceedings of the 4th conference on* (2007), vol. 7, pp. 382–85.

[10] Ojeda, L., and Borenstein, J. Personal Dead-reckoning System for GPS-denied Environments. In *Safety, Security and Rescue Robotics, Proceedings of the IEEE International Workshop on* (2007).

[11] Perry, J. *Gait Analysis - Normal and Pathological Function.* SLACK Incorporated, 1992.

[12] Roetenberg, D., Luinge, H., and Slycke, P. *Xsens MVN: Full 6DOF Human Motion Tracking Using Miniature Inertial Sensors.* Xsens Technologies B.V., April 2009.

[13] Welch, G., and Foxlin, E. Motion tracking: no silver bullet, but a respectable arsenal. *Computer Graphics and Applications, IEEE 22*, 6 (2002), 24–38.

[14] Yeh, S.-Y., Chang, K.-h., Wu, C.-I., Chu, H.-H., and Hsu, J. Y.-J. Geta sandals: a footstep location tracking system. In *Location- and Context-Awareness* (2005), vol. 3479/2005 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, pp. 120–131.

[15] Young, A., Ling, M., and Arvind, D. K. Orient-2: A realtime wireless posture tracking system using local orientation estimation. In *Embedded networked sensors, Proceedings of the 4th workshop on* (2007), pp. 53–57.

[16] Young, A. D., Ling, M. J., and Arvind, D. K. Distributed estimation of linear acceleration for improved accuracy in wireless inertial motion capture. In *Information Processing in Sensor Networks, Proceedings of the 9th The International Conference on* (2010), p. In press.

[17] Yun, X., Bachmann, E. R., Moore, H., and Calusdian, J. Self-contained position tracking of human movement using small inertial/magnetic sensor modules. *Robotics and Automation, Proceedings of the IEEE International Conference on* (2007), 2526–2533.