An Experimental Study of Wireless Connectivity and Routing in Ad Hoc Sensor Networks for Real-Time Soccer Player Monitoring[☆]

Vijay Sivaraman[a,*], Ashay Dhamdhere[a], Hao Chen[a], Alex Kurusingal[a], Sarthak Grover[a]

[a]*School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, NSW, Australia 2052*

Abstract

Live physiological monitoring of soccer players during sporting events can help maximize athlete performance while preventing injury, and enable new applications for referee-assist and enhanced television broadcast services. However, the harsh operating conditions in the soccer field pose several challenges: (a) body-mounted wireless sensor devices have limited radio range, (b) playing area is large, necessitating multi-hop transmission, (c) wireless connectivity is dynamic due to extreme mobility, and (d) data forwarding has to operate within tight delay/energy constraints. In this paper, we take a first step towards characterising wireless connectivity in the soccer field by undertaking experimental work with local soccer clubs, and assess the feasibility of real-time athlete monitoring. We make three specific contributions: (1) We develop an empirical profile of radio signal strength in an open soccer field taking into account distance and body orientation of the athlete. (2) Using data from several soccer games we profile key characteristics of wireless connectivity, highlighting aspects such as small power-law inter-encounters and link correlations. (3) We develop practical multi-hop routing algorithms that can be tuned to achieve the right balance between the competing objectives of resource consumption and data extraction delay. We believe our study is the first to characterise the wireless environment for mobile sensor networks in field sports, and paves the way towards realisation of real-time athlete monitoring systems.

Keywords: Athlete Monitoring, Body-Worn Devices, Dynamic Networks

[☆]This is an extended version of our papers presented at IEEE LCN'10 [20], IEEE WiMob'10 [24], and IEEE SensApp'10 [35].

[*]Corresponding author. Address: School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, NSW, Australia 2052. Phone: +61 2 9385 6577. Fax: +61 2 9385 5993. Email: vijay@unsw.edu.au

Email addresses: vijay@unsw.edu.au (Vijay Sivaraman), ashaydhamdhere@gmail.com (Ashay Dhamdhere), music.tennis.boxer@gmail.com (Hao Chen), alexkurusingal@gmail.com (Alex Kurusingal), shahifaqeer@gmail.com (Sarthak Grover)

1. Introduction

Advances in sensing and communications technologies are enabling new low-cost and lightweight devices that allow measurement and remote monitoring of an individual's vital physiological signs such as ECG, temperature and oxygen saturation levels. Such technology, though designed primarily for the healthcare industry, is being adapted to the massively popular and growing field of sports science, specifically for the purpose of athlete monitoring.

Biomedical technology has long been used by professional coaches and trainers in striving to push their athletes' bodies to the edge of its capabilities. However, much of this examination of the body has been performed under laboratory conditions where results attained in the artifical environment may not parallel those observed in competition [1]. Devices are now starting to emerge in the market that are making the leap from monitoring athletes in training (e.g. SPI Elite [2] platform from GPSports) to monitoring them during competition (e.g. e-AR [3], VxLog [4], and WiMu [5]). We are partnering with Toumaz Technology Ltd. in the UK who are manufacturing a platform called Sensium^{TM}[6] that integrates low power wireless technology with miniaturised sensors and lightweight flexible batteries [7]. This platform, weighing under 10 grams, will allow non-intrusive collection and real-time wireless transmission of athlete physiological data during competition.

We seek to apply the above wearable platforms to monitoring athletes in field sports, specifically soccer. Soccer is a hugely popular sport throughout the world, and attracts large financial investment, particularly in Europe. Several soccer clubs in the UK have expressed great interest in monitoring their athletes on the field, predominantly to reduce the risk of injury and improve player substitution decisions. Soccer organisers have also expressed interest in using real-time position and impact information for referee-assist services, and television channels are eager to augment live broadcasts with player parameters (e.g. heart-rate during clutch events, speed and acceleration, impact levels during collisions, etc.) so as to heighten the level of engagement for audiences.

While hardware platforms for athlete monitoring are maturing rapidly, there is much research needed in developing communication protocols that can operate under the unique conditions arising in the soccer field: (a) Rapid acceleration and impact are part of the sport, and this restricts the monitoring device to be small, lightweight, unobtrusive and non-protruding so that the players' degree of freedom is not limited. This is in contrast to devices tried in sports such as rowing [8] or cross country skiing [9] that have form-factor akin to a mobile phone. Consequently, monitoring devices for soccer can be expected to have extremely limited battery power and restricted radio range, placing severe energy and reach constraints on the communication protocols. (b) The playing area in soccer is very large at over $4000m^2$. Given the limited transmission range of body-worn devices, coupled with attenuation effects arising from attachment to the human body (profiled later in this paper), real-time extraction of player

data would require multi-hop routing. One-hop communication from the device to base-station, such as proposed for ice-hockey in [10], or the protocols proposed in [11] for monitoring team-sports such as basketball and volleyball having a small playing area, would not suffice for soccer. (c) Soccer players move very rapidly in the field, and this makes the topology highly dynamic at short time-scales (seconds). Designing routing mechanisms that can deliver data to base-stations within stringent time and energy constraints over multiple hops in this time-varying environment promises to be challenging.

To the best of our knowledge there is neither any prior work in characterising the operating environment for a mobile body-worn wireless sensor network in a soccer field, nor are there protocols in the literature that are suitable for real-time athlete data extraction in such an environment. In this paper we undertake experimental work in which we outfit several soccer clubs with sensor devices and collect data on their movement and connectivity (with each other and with base-stations around the field) over multiple games. Using the collected data we make three contributions:

- Our **first** contribution develops an empirical model of the strength of the radio signal emanating from an athlete's body-worn device. Unlike previous disk models that assume isotropic propagation, we show that the signal strength varies with both angular orientation and radial distance when the transmitter is worn against the body. Using our empirical data, we derive an analytical fit that provides accurate characterisation of the radio reach of an athlete's sensor device. This characterisation will allow us to generate the wireless topologies arising in the soccer field using empirical data that tracks the location of players during the game.

- Our **second** contribution uses empirical data to provide a stochastic characterisation of key aspects of the dynamic wireless topologies arising during a soccer game, such as the number of wireless neighbours of a player (indicating the number of alternate routes that may be available), distributions of the encounter and inter-encounter times between players (indicating the length of time for which routes may persist or vanish), and correlations that exist amongst links (i.e. the presence of link between one pair of players can positively or negatively affect the probability of link between another pair). Using this characterisation, we develop a novel mathematical model for generating dynamic topologies representative of real soccer games, that accurately depict the auto-correlation and cross-correlation structure of links, from which derived metrics such as inter-contact times and neighbourhood distribution follow. We believe our model is the first in the literature to directly generate connectivity topologies for arbitrarily specified link (auto- and cross-) correlations, which has general applicability beyond field sports.

- Our **third** contribution establishes, using connectivity traces measured in real games as well as connectivity data inferred from player location, that multi-hop routing has the potential to substantially reduce athlete data

Figure 1: Equipment used in the experiments: Motes and GPS units for each player, and 8 base-stations with high-gain antennas

extraction delays compared to single-hop transmission from the body-worn device to a base-station. We then develop a class of practical multi-hop routing algorithms that can be tuned to achieve the desired trade-off between energy performance and delay, showing that real-time monitoring of athletes using a mobile wireless sensor network of ultra-light-weight body-wearable devices is not a distant reality but very feasible in the near future.

The rest of this paper is organised as follows: In Section 2 we describe our experimental setup. Section 3 presents empirical data and a mathematical model of radio signal strength from a sensor device mounted on an athlete's body. Using this model we perform, in Section 4, an extensive characterisation of the dynamics of the wireless topology arising in soccer games, and develop an analytical model that can generate such topologies. In Section 5 we show that multi-hop routing is required for real-time data delivery, and develop practical multi-hop routing schemes that can trade-off energy performance for delay. Section 6 concludes our work and presents directions for future research.

2. Experimental Setup

With the objective of gaining an understanding of wireless connectivity during soccer games, we procured equipment including wireless sensor devices, GPS units, and base-stations with high-gain antennas, as shown in Fig. 1. We then

(a) Mote on right arm (b) GPS on left arm

Figure 2: Mounting of MicaZ mote and GPS unit on player

outfitted soccer club players with wireless monitoring devices and/or GPS devices, as shown in Fig. 2, and collected data over six competitive league games spanning two seasons (2008-2009 and 2009-2010). Animations from three of these games can be seen on our project web-page [12]. We outfitted players of the University of New South Wales Football Club (UNSWFC) first-division men's team for four games, and players of the Putney Rangers sixth-division team for two games (in one of the latter games we were able to outfit players from the opposing team as well). We wish to state here that getting perfect data (from every player for the entire game) turned out to be very challenging, due to the frequent impacts during the game that dislodged or damaged some of the devices. Nevertheless, the data we were able to gather gives us a sufficiently good picture of how wireless connectivity evolves during the soccer game. Space constraints prevent us from discussing all the individual games we monitored, so in the rest of this paper we will focus on two specific games (for reasons outlined next) when discussing specific measures, and generalise our observations to include the other games where possible.

2.1. Game 1 (Feb 2009): Connectivity Data

The first game we focus on was played in February 2009 by the first-division UNSW Football Club. The game was played on a full size field with dimensions 93m x 70m. Each of the 11 players wore a monitoring device on their arm, and 8 base-stations were positioned (at a height of about 1m from the ground) along the sidelines of the playing area. We note that for this game the base-stations used their standard quarter-wavelength dipole antenna (for subsequent games we procured a bigger high-gain antenna). Fig. 3 shows the nominal playing positions and associated node identification numbers. Unfortunately the devices worn by players 2 (back) and 4 (left back) were damaged during play and we could not obtain data from them, as was base-station B3 which got hit by the ball.

The body-worn devices we used for this game were the MicaZ motes [13] from Crossbow technologies. These are off-the-shelf devices operating in the 2.4GHz

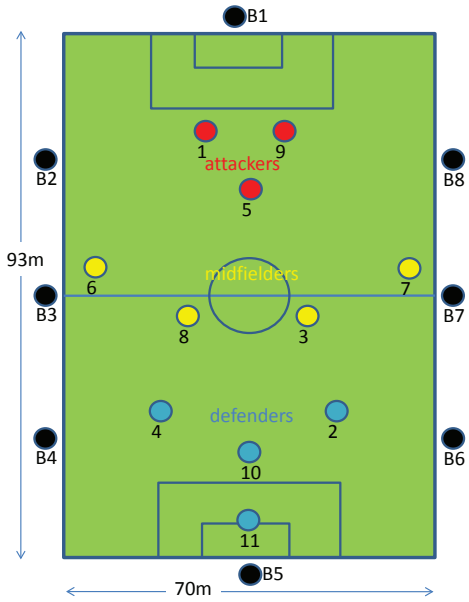

Player Number	Playing Position
1	Center Attack
2	Back
3	Centre Midfield A
4	Left Back
5	Center Forward
6	Left Wing
7	Right Wing
8	Centre Midfield B
9	Striker
10	Center Back
11	Goalkeeper

Figure 3: Player default positions (Game 1)

band that are readily available today. Though they were not designed for body-worn applications, they have been used before for body health monitoring, such as in Harvard's Code Blue project [14], and in our own prior work [15] in profiling the body channel for patients with chronic illnesses. We intend to replace these with emerging platforms custom-built for body-area-networking as they become available.

One of the foremost challenges we faced was in finding a good way to mount sensor nodes on athletes. Taking into consideration aspects such as attenuation of the wireless signal by the body, ease and stability of attachment, and possibility of damage to the device itself, we decided to go with an arm mounted attachment using an arm-band. We also tried other mounting positions (e.g. back), but found such mountings to either cast a larger "shadow" region of poor signal, or create discomfort for the athlete due to clothing impediments or increased chance of injury/damage during a fall. We therefore proceeded with an arm-mounted position for all our subsequent studies. A detailed study of the wireless propagation from the body-worn device will be presented in the next section.

We implemented software on each of the body-worn devices such that it broadcasts, once every second, at the highest available power level of 1mW (0dBm), a packet containing its unique identifier and a sequence number. All devices (body-worn as well as base-stations) that successfully receive this packet record this event in their on-board memory. As the game proceeds, each

node (and base-station) will be cataloguing which other nodes it could hear at each time instant. To prevent collisions in-the-air, each second is divided into 11 slots each of approximate duration 90ms, and each of the 11 body-worn devices is given a unique such slot for transmission every second. Just prior to commencement of the game, the master base-station sends a clock synchronisation message to all nodes, upon receipt of which each node starts recording connectivity data in on-board memory. Data collection stops after 25 minutes (due to limitation on on-board space for storing the connectivity data), and at the end of the game data from each node is extracted by the master base-station for off-line analysis. A Java animation of the connectivity we observed for this game can be seen at the project web-page `http://www2.ee.unsw.edu.au/~vijay/athlete/game2009feb`. Note that the players are static (at their default playing position) in this animation (which only shows how connectivity changes), since we do not have location information for the players. We recommend the reader to view the animation from 900s to 1500s, since we use data from that 10-minute interval for our analysis as there were no substitutions and no play stoppages during that period.

2.2. Game 2 (Aug 2010): Location Data

The second game we focus on was played by the sixth-division Putney Rangers team in August 2010. In this game we took a different approach and outfitted all players from both teams with GPS tracking devices (in addition to wireless sensor devices) for two reasons: (a) By knowing the location of all players at all times, we have flexibility in reconstructing radio connectivity for differing radio characteristics found in emerging body-wearable devices. For example, the Sensium^{TM}[6] has a maximum transmit power of -6dBm, which is lower than the MicaZ mote's default transmit power of 0dBm; therefore, using the Sensium^{TM}as the body-worn device would result in sparser wireless connectivity than with the MicaZ mote. By having location information, the wireless connectivity can be generated for given radio transmit strength, allowing flexibility in study of real-time athlete monitoring for different device characteristics. (b) The GPS units (BT-Q1300ST GPS sports recorder from QStarz [16]) are smaller, lighter, and have better attachment (via a supplied arm-band) than the wireless sensor nodes, so we did not lose any of the location data in this game from device damage, unlike the more bulky MicaZ motes that we lost data from for several players.

We have developed a tool that animates the player locations to show how the game evolves, and super-imposes on it the wireless inter-connectivity amongst players (computed using the model derived in the next section). A snapshot of the tool is shown Fig. 4, and a web-version of the animation can be viewed at the project web-page `http://www2.ee.unsw.edu.au/~vijay/athlete/game2010aug`. Note that location sampling was done every second but we speed up the animation by a factor of 10 for ease of viewing (the speedup can be adjusted by the viewer). Also note that due to substitutions not all players were playing at all times. Unfortunately we do not have a way of tracking the ball during the game.

Figure 4: Animation tool snapshot showing player locations and inter-connectivity (Game 2)

Knowing the second-by-second location of all players during the game allows us to contruct the wireless interconnectivity topology between players (and to base-stations around the field) for given assumption on radio transmit power from the body-worn device. This in turn lets us study the feasibility of real-time extraction of athlete's vital physiological signs during the game as a function of device characteristics. Before we can do that, we need a realistic model of signal strength propagation from a device worn by an athlete, which is the topic of the next section.

3. Modeling Radio Propogation around the Athlete's Body

The objective in this section is to develop a model of radio signal propagation from a sensor device worn by an athlete, which can then be used to deduce inter-connectivity amongst soccer players during the game, based on their location and orientation. Unlike much of the prior work which assumes symmetric (disk-shaped) radio signal strength in all directions from the source, we observed in our experiments that the signal strength for body-worn devices is heavily influenced by body orientation, due to absorption by the body.

There has been some prior work in evaluating the influence of the human body on wireless signal propagation characteristics. For instance, [17] studies the effect of the human body on WiFi propagation from portable computers, and shows that there is a 25dB loss in signal strength when the human body is

in the way. The propagation of radio signals in and around the body has been studied extensively by Hall and Yao [18, 19]. Their study concentrates mostly on on-body networks, where the transmitter and receiver are both on the same body, and shows that absorption due to the water content of the human body resulted in a 40dB path loss. We have not been able to find a characterisation in the literature of radio signal strengths between a body-mounted transmitter and an off-body receiver, and in what follows we develop such a model based on empirical data, building on our recent work in [20].

3.1. Experimental Set-Up

We used MicaZ motes from Crossbow Technologies [21] running the TinyOS operating system as our transmitter and receiver. As the transmitter represented the body-worn sensor device, it used a 1/4 wavelength dipole antenna that comes standard with the Mica motes. The receiver represented a base station (perhaps located on the periphery of the soccer field in an athlete monitoring application), and therefore used a bigger high-gain (+12dBi) antenna from TP-Link [22], mounted on a tripod.

Our objective was to characterize how the received signal strength varied with the relative position between the transmitter and receiver. The transmitter sent packets at a fixed rate of 4 packets per second, at a fixed power level of 1mW (0dBm). Upon successfully receiving a packet, the receiver computed the Received Signal Strength Indicator (RSSI) of the received packet and sent this value to a laptop computer over the serial port. In the micaZ motes the RSSI value is an 8-bit number obtained by sampling the onboard ADC during packet reception. The RSSI value was then converted to a dBm value by subtracting 45 [23]. Our experiments were performed in an open soccer field away from any sources of interference.

3.2. Experimental Observations

We performed two experiments: in the first experiment, our objective was to ascertain the propagation pattern in free space, so as to have a baseline against which to compare the effect of the human body. The transmitter was mounted at the top of a non-conducting pole at a height of 1.5m above ground level, while the receiver remained stationary. We then increased the distance between transmitter and receiver in steps of 1m, and recorded the RSSI reading at each step. This was repeated until reliable reception could not be obtained (where reception was considered *reliable* when there was no packet loss). We verified that the signal strength was roughly isotropic (i.e. identical in all directions), and show in Fig. 5 the 3-D contour plot of the measured signal strength highlighting that the RSSI depends only on distance and is not dependent on the angle between the transmitter and receiver.

In the second experiment the transmitter was mounted on the right arm of a test subject, as shown in Fig. 2(a). The subject rotated his body in 15 degree increments from $0^\circ \rightarrow 345^\circ$ with respect to the receiver. The observed signal strength measurements are shown as a 3D plot in Fig. 6(a). As expected,

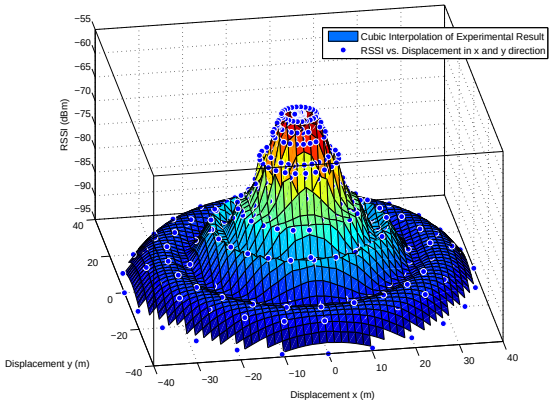

Figure 5: Interpolated surface fit of Free Space Experimental Data

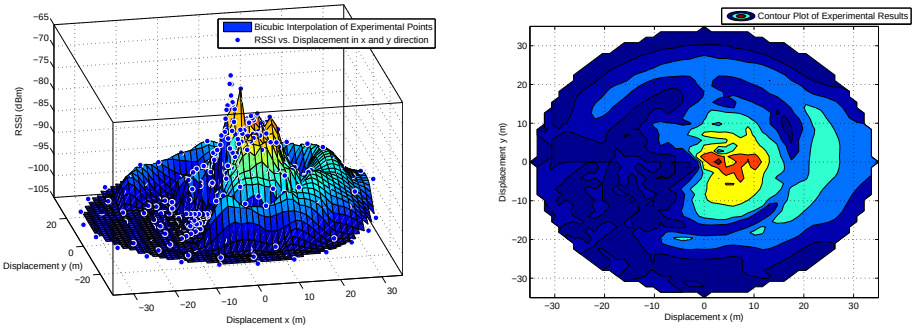

(a) 3D plot

(b) 2D projection

Figure 6: Signal strength around athlete's body (experimental data)

the RSSI contours now show very significant reduction in signal strength when the body is in the way. When the body is not blocking the signal, we observe 30m of uninterrupted range. On the other side, however, the range is only 2m before reception drops completely. The highest recorded RSSI occurred at an orientation of approximately 45°, at a distance of about 1m.

3.3. Analytical Model

Using the empirical data, we now derive an analytical model to deduce the signal strength as a function of distance from the transmitter and the orientation relative to direct line-of-sight. Free-space models typically set the received signal strength to fall with distance r as a power-law: $RSSI = a * r^{-b}$ where b is close to 2.5 and a is a proportionality constant incorporating effects of the antenna, transmit power and environmental variables. We chose to keep our formulation

consistent with this notation, where the a and b terms are now parameterises by angle θ at which the receiver is orientated to the direct line-of-sight with the body, and the RSSI for the body-worn scenario hence takes the form:

$$RSSI = a_\theta * r^{-b_\theta} \tag{1}$$

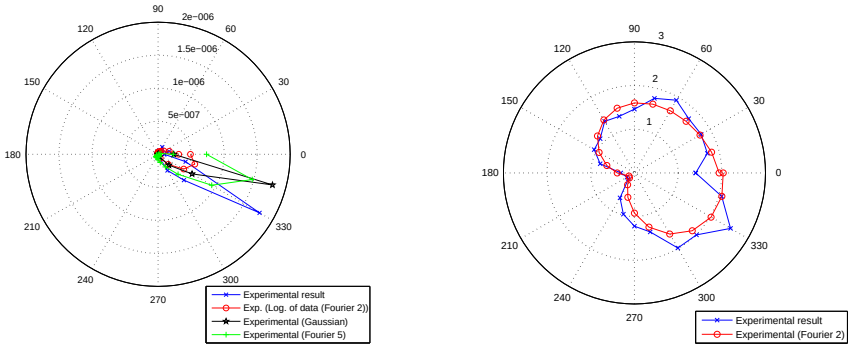

(a) a_θ and associated fits

(b) b_θ and associated fits

Figure 7: Polar plots of a_θ and b_θ and associated fits

The best-fit values of a_θ and b_θ for the various values of θ (at 15° increments) obtained from the experimental data are shown as polar plots in Fig. 7(a) and 7(b) respectively. The figures also show anlytical fits to a_θ and b_θ from three models: Fourier, Sum of Sines, and Gaussian. We evaluated each fitting method on the basis of simplicity of the expression and goodness of the fit.

The power-law exponent b_θ was fitted best by a second order Fourier series of the following form (with $R^2 \approx 0.9033$):

$$\begin{aligned} b_\theta \quad = \quad & 1.27 + 0.8086 * \cos(\theta * 0.9726) + 0.1851 * \sin(\theta * 0.9726) \\ & -0.1396 * \cos(2 * \theta * 0.9726) - 0.3049 * \sin(2 * \theta * 0.9726) \end{aligned} \tag{2}$$

For parameter a_θ it was found that fitting $\log_{10} a_\theta$ resulted in a better goodness of fit ($R^2 \approx 0.9603$) for the same complexity of expression. The best fit was obtained for a second order Fourier series of the following form:

$$\begin{aligned} -\log a_\theta \quad = \quad & 7.868 - 1.551 * \cos(\theta * 0.9893) \\ & -0.1774 * \sin(\theta * 0.9893) + 0.1882 * \cos(2 * \theta * 0.9893) \\ & +0.5404 * \sin(2 * \theta * 0.9893) \end{aligned} \tag{3}$$

3.4. Validating the Model

We now validate the model by comparing its estimates with the experimental data. Fig. 8(a) shows a 3D surface which was interpolated using data points generated by substituting (3) and (2) into (1), while the dots in the figure represent the experimental data points, and the match is found to be quite good. In Fig. 8(b) we show a contour projection of the 3D surface, and comparing this with Fig. 6(b), we see that the model predicts close to 40m of uninterrupted coverage in the right hand direction, whereas the experimental result suggests that a little over 30m range is possible.

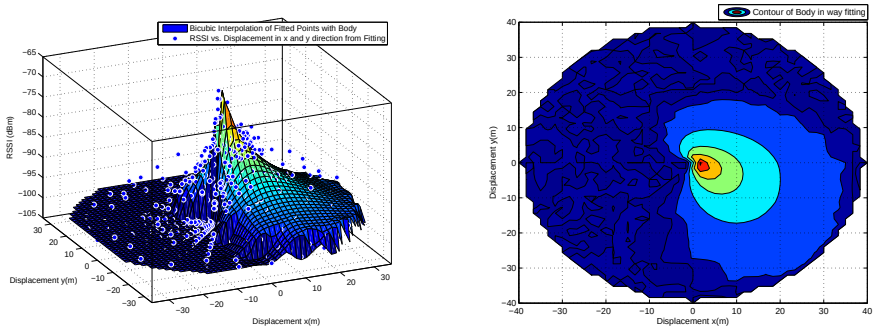

(a) 3D plot (b) 2D projection

Figure 8: Signal strength around athlete's body (from model)

3.5. Discussion

It is known that the human body absorbs radio signals at 2.4GHz. As a result, we expect very high attenuation for orientations where the body significantly shadows the signal (it has been further shown that at these orientations the dominant component of the signal arises from creeping waves around the body [19]). It should be noted that we are not measuring the signal directly on the body, however, the above results are still significant especially at orientations between 150 and 270 degrees.

Noting further that the a_θ values denote the received signals at a distance of 1m from the transmitter, along a given orientation, we see that a_{180} ($= 10^{-9.646}$) is two orders of magnitude below a_0 ($= 10^{-7.028}$); this loss is consistent with the loss experienced by a creeping wave traveling halfway around the human body. Similarly, we compare the b values at different orientations. It is found that at orientations between 90 and 270 degrees (where the wireless range is significant), the b values lie in the range -1.2 to -2. While this is a higher value than that for free space (leading to a slower decay), it is combined with a much lower value of a, as compared to free space.

Our analytical model that estimates signal strength as a function of distance and orientation of the receiver from the transmitter provides a realistic mechanism for deriving connectivity between body-worn devices, something that has not been reported in the literature before. In subsequent sections we will show that our mechanism helps us characterise the wireless connectivity amongst soccer players during a game using their location information, allowing us to develop efficient protocols for real-time extraction of athlete data.

4. Profiling Dynamic Wireless Topologies in a Soccer Game

With a view towards designing communication protocols that are suited to the soccer field, in this section, we profile the dynamics of the wireless topologies that arise in the soccer field. To the best of our knowledge no data or

characterisation of real soccer games is available in the public literature today (apart from our own preliminary characterisation in [24]), and our data, viewable on our project web-page `http://www2.ee.unsw.edu.au/~vijay/athlete`, can serve as useful input for other researchers studying field-sports.

Using data on connectivity (from body-worn MicaZ motes in Game 1) and position (from GPS units worn in Game 2) from the soccer games, in this section we make two new contributions: First, we provide a **stochastic characterisation** of key aspects of the topology, such as the number of wireless neighbours of a player (indicating the number of alternate routes that may be available), distribution of the flight length and inter-encounter times between players (indicating the length of time for which routes may persist or vanish). Additionally, we show that several links exhibit strong correlations with each other, i.e. the presence of link between one pair of players can affect the probability (positively or negatively) of link between another pair. Second, we propose a novel **mathematical model** for generating dynamic topologies that stochastically match empirical traces. Our model explicitly considers the underlying auto-correlation and cross-correlation structure of links, from which derived metrics such as inter-contact times and neighbourhood distribution follow. Our model is unique in being able to directly generate connectivity topology for arbitrarily specified link (auto- and cross-) correlations, with potential application to areas beyond field sports. Our study in this section sets the stage for developing new routing mechanisms suited to soccer player monitoring in the next section.

4.1. Profiling Player Connectivity

Recall that we have second-to-second connectivity between (most of the) player-worn devices from Game 1. For Game 2, we infer this connectivity (for all players) by using the second-by-second GPS location of the players as follows: for each time-step, we compute distance between each player and base-station (or another player) using their positions, as well as their relative orientation (from direction of movement deduced from change in position). We then apply our propagation model of the previous section to compute the signal strength in each direction; if the received signal strength is computed to be higher or equal to the receiver sensitivity (set to -100dBm), the link is present at that time-step, otherwise it is absent. Using the connectivity data thus obtained, we now profile various aspects of player connectivity, such as the number of neighbours and inter-encounter durations, as well as correlations amongst links. Though we recognise that each soccer game is different and data acquired from repeat trials would undoubtedly yield a different composition of results, our aim is to highlight key common characteristics and trends associated with player connectivity arising in real soccer games.

4.1.1. Number of Neighbours

In Fig. 9 we show the probability distribution of the number of neighbours for selected nodes (i.e. number of nodes whose transmission can be heard by the selected node). To give a flavour of the diversity we pick players from forward

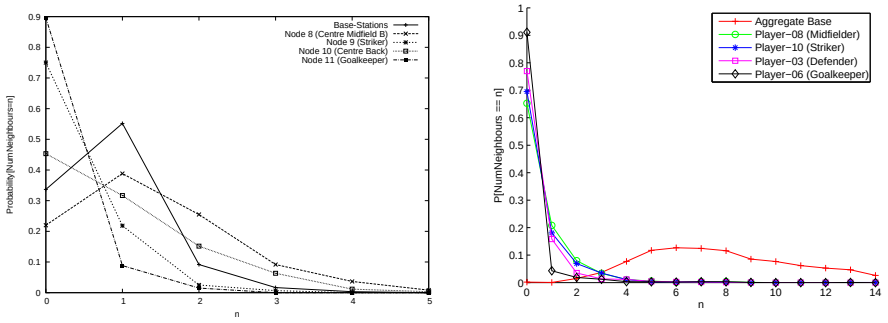

(a) Game 1 (b) Game 2

Figure 9: Neighbour distribution for base-stations and representative players in striker, midfielder, and defender positions

(striker), middle (centre midfielder), and backward (defender or goalkeeper) playing positions, as well as the base-stations (aggregated together). Some interesting observations can be made from the figure:

- In general the average number of neighbours for a player is quite low, predominantly due to small range of the body-worn device compared to the playing area of the soccer field. In Game 1 the number of transmitters within range of base-stations (aggregated) is also quite low (we did not use high-gain antennas for the base-stations in this game), but this improves significantly in Game 2 wherein we used high-gain antennas on the base-stations. As we will show in the next section, improved connectivity to the base-stations reduces delay in extraction of athlete data via single-hop transmission, but not as much for players in the centre of the field, who benefit from multi-hopping through other players.

- Connectivity to other players varies with playing position: for example, in both games it can be seen that midfielders have better connectivity, due to central location on the field, than the striker or goalkeeper who are more likely to be at the extremes of the field. This information can be exploited by routing algorithms.

4.1.2. Athlete Flight-Lengths and Speeds

We take a look at two aspects of player mobility that have a bearing on their connectivity: flight-path length (i.e. distance for which a player moves in roughly the same direction) and its relation to the athlete's speed. There has been recent evidence that humans follow a Levy Walk, namely the distance people travel in roughly the same direction before pausing or changing directions has a power-law distribution [25]. We seek to verify if this holds in a soccer game. Using the second-by-second GPS location information we gathered in Game

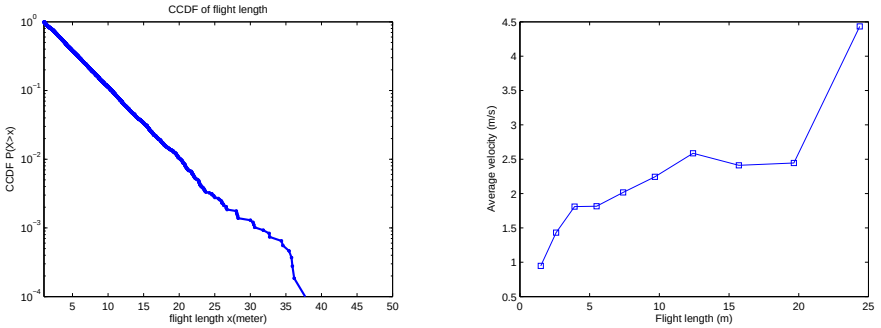

(a) Flight length

(b) Flight length and Speed

Figure 10: Distribution of flight-path length and its correlation with speed (Game 2)

2, we apply the rectangular box algorithm from [25] to determine the flight-length, and plot its complementary cumulative distribution function (CCDF) in Fig. 10(a). We find that the CCDF is nearly a straight line on log-linear scale, implying that flight-path lengths are exponential rather than power-law in a soccer game. Also, in Fig. 10(b) we plot the athlete's speed (obtained from change in athlete's GPS location over each second) during the corresponding flight-path, and find that there is a monotonic relationship, i.e. longer flight are associated with greater velocity. Though we do not directly use flight-path length and speed in this work, they are interesting metrics providing evidence that movement on the soccer field has different fundamental characteristics to regular human movement (which has been modeled by many authors in the literature), and requires new models like the one we present later in this section.

4.1.3. Inter-Encounter Times

Another metric that is known to have an important bearing on the route-selection algorithm in mobile ad-hoc networks is the inter-encounter time [26] (also known as inter-meeting or inter-contact time) between nodes. Most prior studies have relied on exponentially distributed inter-encounter times for tractable anaysis of routing performance; however, recent studies such as [27] have shown that non-exponential behaviour can lead to unbounded routing delays. To see which model best fits the soccer field environment, in Fig. 11 we show the distribution of the inter-encounter time amongst all pairs of nodes, as well as between all transmitters and a specific receiver (midfielder, chosen for its rich connectivity), for both games.

We found that the Complementary Cumulative Distribution Function (CCDF) of the inter-encounter time did not follow an exponential distribution, and therefore we depict the inter-encounter time CCDF on log-log scale in the figure. We find that the head of the inter-encounter delay curve is roughly linear on log-log scale, indicative of power-law behaviour in that range. The power-law exponent

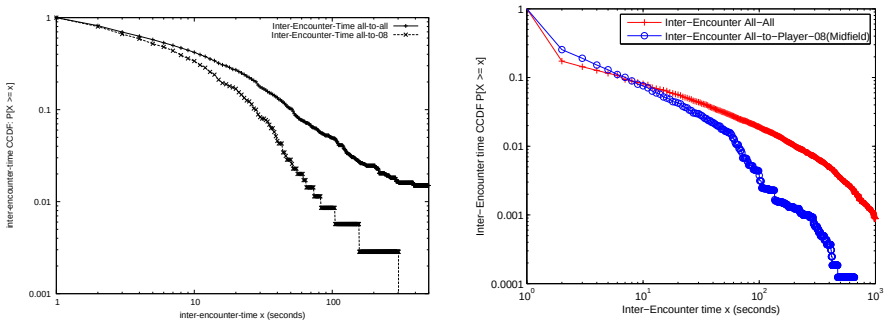

(a) Game 1 (b) Game 2

Figure 11: Distribution of inter-encounter times for Game 1 and Game 2 on log-log scale

in this region is estimated at around $\alpha \approx 1.6$. Though [27] estimates analytically that $\alpha < 2$ leads to unbounded routing delays, it does so by extrapolating the inter-encounter delay tail to infinity as a power-law. Our experimental data shows that the tail of the curve flattens out (on log-linear scale), and in this region inter-encounters are better modelled as exponential. This combination of power-law and exponential behaviour is consistent with reported mixtures [28] seen in inter-meeting times for regular human activity, and result in bounded routing delays unlike the pessimistic estimates in [27].

To characterise the encounter and inter-encounter distributions and their auto-correlations in a succinct way (which we employ later in this section for our model), we borrow a technique used for the analysis of long-range dependent (LRD) traffic. Considering a link between a pair of nodes, at a given time step, we use a 1 to depict presence of the link and 0 its absence. For this link, we therefore have from our experimental data a time-sequence of 0s and 1s. We consider this sequence in blocks of 2^s samples, for given s, and for this resulting sequence we compute the mean, variance, and coefficient-of-variation $\beta(s)$ (in effect these metrics are computed at time-scale 2^s). Log-log plots of $\beta(s)$ versus s are routinely used in the literature to depict self-similarity and to estimate the corresponding Hurst parameter $H \in [0.5, 1)$. In Fig. 12 we show such a plot for several links (we picked two links each from centre, forward, and backward playing positions), and observe that the curves can be approximated as straight-lines with slope $-(1 - H)$, yielding a Hurst parameter $H \approx 0.75$. This single-parameter captures in a succinct way the link auto-correlations, and will be used in the connectivity model we develop later in this section.

4.1.4. Link Correlations

Unlike many mobile ad-hoc networks in which we can reasonably assume that users move independently, in a soccer game we would expect player movements to have significant correlations. For example, when the team is attacking the

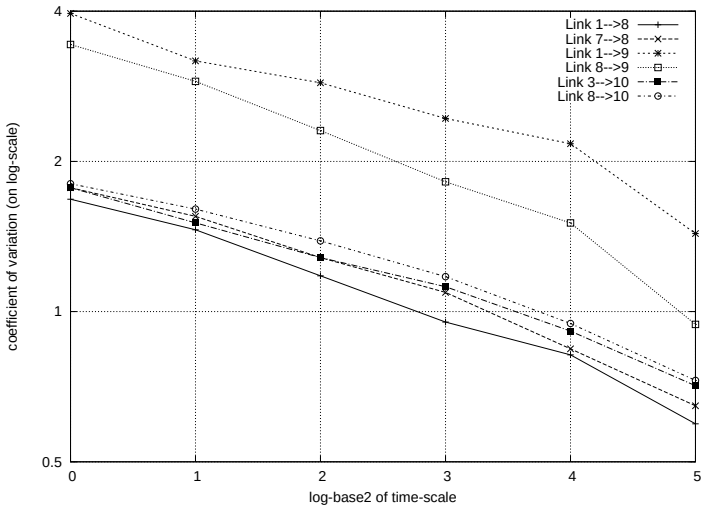

Figure 12: Log-log plot of link coefficient-of-variation versus time-scale (seconds)

opponent's goal, several players in the forward and midfield positions can be expected to move towards the opponent's goal simultaneously, and conversely when the home goal is being attacked the defenders and midfielders will likely fall back towards the home goal to protect it. This leads to correlations amongst links, an aspect which can have a significant impact on the performance of routing algorithms.

Correlations are computed as follows: if x_t is a binary variable that is 1 or 0 depending on whether link x is present or absent at time step t, then the cross-correlation at time lag k between two links x and y is given by [29, Sec 12.1.2]:

$$\rho_{xy}(k) = \frac{\frac{1}{n}\sum_{t=1}^{n-k}(x_t - \bar{x})(y_{t+k} - \bar{y})}{\sigma_x \sigma_y}, \quad k = 0, \pm 1, \pm 2, \ldots \tag{4}$$

where n is the number of sample points, $\bar{x}$ is the estimated mean and σ_x the estimated standard deviation of x.

In Fig. 13(a) we depict some correlations from Game 1: specifically between node 3's (centre midfield A) and node 10's (centre back) links to node 8 (the centre midfield B) for lags in the range $[-20, 20]$ seconds. Two things are noteworthy from this plot: (a) the correlations are positive, meaning that when node 3 is close to node 8, node 10 is also likely to be close to node 8; this suggests nodes 3 and 10 move in a co-ordinated way quite often, and (b) the correlations are high (> 0.2) for lag close to 0, and decay rapidly as the lag moves away from 0. This is not surprising, because the fast nature of the game implies that the locations of the players can vary significantly from one minute to the next, making them nearly independent. In Fig. 13(b) we show the correlation between node

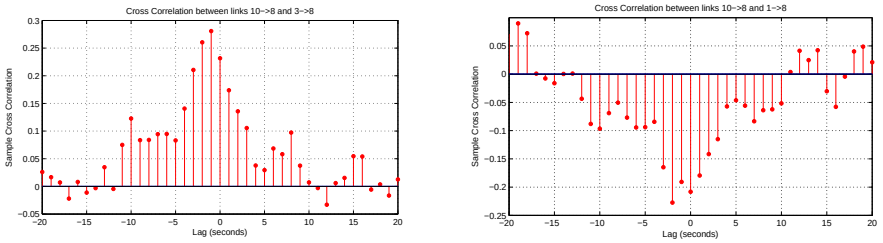

(a) correlation between $3 \to 8$ and $10 \to 8$ (b) correlations between $1 \to 8$ and $10 \to 8$

Figure 13: Correlation of (a) link $3 \to 8$ with $10 \to 8$ and (b) link $1 \to 8$ with $10 \to 8$, as a function of lag (in seconds) from Game 1

1's (centre attack) and node 10's (centre back) with node 8 (centre midfield B). This time we notice that the correlations are predominantly negative (< -0.2 for lags close to 0), which is understandable: when the team is attacking, the midfielder is more likely to be close to the striker and far from the defender, while the converse is true when the team is defending their own goal. Again we notice that the anti-correlations decay with time due to the rapid movement of players in the game.

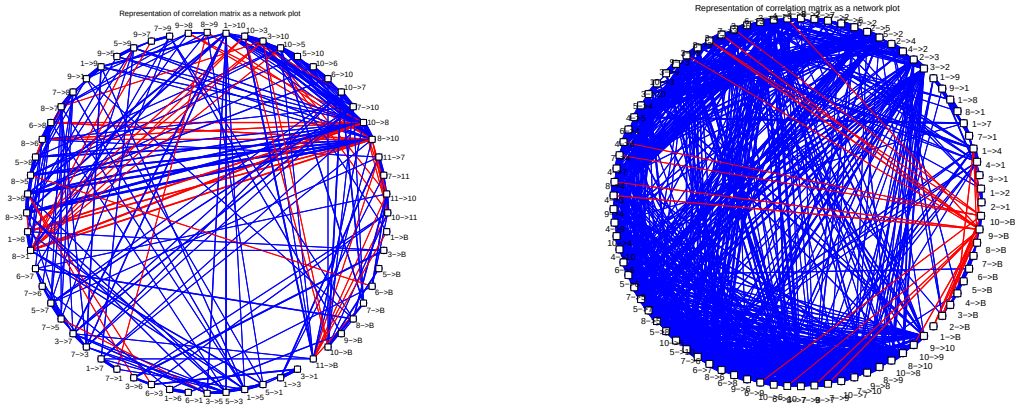

(a) Game 1 (b) Game 2

Figure 14: Correlations for Game 1 and Game 2 depicted via inter-connections

Having seen specific examples of correlated and anti-correlated links, we now depict observed correlations amongst all pairs of links in Fig. 14. We place all links as nodes on a circle in the figure, and draw a line between two nodes if the corresponding links have significant correlation: blue lines depict positive correlation while red lines depict negative correlation, and the higher the correlation (or anti-correlation), the thicker the line. Also, links have been

ordered on the circle so that the two directions of the link are adjacent to each other (so that correlations between the two directions of a link do not clutter the plot). To eliminate random chance of correlated values, we also estimate the P-value [30] (used for statistical hypothesis testing) for each pair, and only retain those that are statistically significant (i.e. have $P \leq 0.05$).

In Game 1 we had only 60 links that had statistically significant ocurrence, while Game 2 has many more links (since we got much more comprehensive data from all players of the team), hence Fig. 14(b) has many more nodes than Fig. 14(a). Nevertheless, we can see that in both games there are significant correlations (both positive and negative) between links, even amongst links that do not have a common endpoint.

4.2. A Connectivity Model for Soccer Players

We now develop a model that can generate synthetic dynamic topologies with similar stochastic properties to those observed empirically. Such a model would be useful in generating long traces to simulate the performance of different routing strategies for soccer player monitoring, and also allow key parameters such as link auto- and cross-correlations (which may depend on a team's playing style) to be varied to study their impact on routing performance.

One approach to modeling wireless connectivity is to model the movement of players on the soccer field. Though mobility has been modeled in the literature for various contexts (see [31, 32] for a survey), ranging from individual node mobility (e.g. Random Waypoint model, Levy Walk model [25], etc.) to group mobility (e.g. Reference Point Group Mobility model and Pursue Mobility model), we found that none of these existing models were a good fit for our data from soccer games. Moreover, we found that modeling mobility of players was difficult unless we could also model mobility of the soccer ball, for which we unfortunately do not have data since we had no way of tracking the ball in our experiments. In this paper we therefore focus on modeling the connectivity (aka wireless topology) between players, rather than inferring connectivity from a mobility model. Only a handful of prior works have directly modeled connectivity: one example is [33] that proposes a statistical encounter-based model in the context of delay tolerant networks (DTNs). However, their model assumes links to be independent, which is inadequate for capturing correlations that we have shown to exist in team sports such as soccer. The model we present next overcomes this important limitation.

4.2.1. Model Requirements

We seek a model that takes the following inputs: (a) Number of players, base-stations, and links, (b) Mean and variance for each link (the link is binary in each time-step: 0 if down and 1 if up), (c) Auto-correlation of the links (to keep the model simple we assume that all links have similar auto-correlations), specified via the Hurst parameter (section 4.1.3) or auto-regressive coefficients (discussed in the next subsection), and (d) Cross-correlation between each pair of links, specified as a covariance matrix.

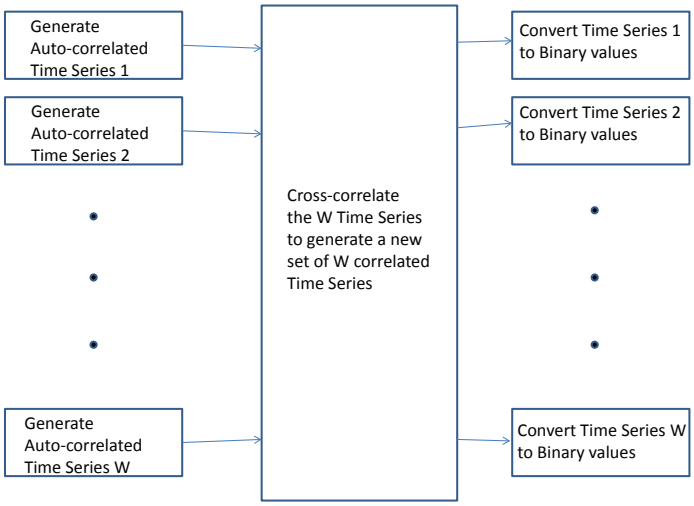

Figure 15: Flow diagram of our model for generating time-varying topology

The model should output for each successive time-step the connectivity topology, i.e., the set of links that are up at that time-step. If an empirical trace is available from which the input parameters were derived, then the generated topology should statistically match the empirical trace in the following metrics: (a) for each link, the on/off (1/0) distribution, (b) the distribution of the number of active links in the network, (c) the distributions of encounter durations and inter-encounter times, and (d) the correlations between every pair of links.

4.2.2. The Model

We use W to denote the total number of links (player-to-player as well as player-to-base). The covariance matrix (which is an input to the model) is denoted by $\mathbf{C}$, and is of dimension $W \times W$. Element C_{ij} denotes the covariance between the i-th and j-th links, and is related to the correlation defined in Eq. (4) by $C_{ij} = \sigma_x \sigma_y \rho_{ij}(0)$ (note that to reduce complexity our model directly incorporates correlation at lag 0 only; correlations at other lags will follow from cross-correlations at lag 0 combined with the auto-correlations of the links). Further, each diagonal entry C_{ii} corresponds to the variance of the binary variable associated with the i-th link. A valid covariance matrix is required to be symmetric (i.e. $C_{ij} = C_{ji}$) and positive definite (i.e. $\mathbf{C} > \mathbf{0}$).

The general flow of our model is shown pictorially in Fig. 15, and broadly consists of three steps: step 1 generates independent random variables, one per link, with the appropriate auto-correlation, step 2 mixes them to create the correct cross-correlations, and step 3 converts them from continuous to discrete

(binary) values so they correspond to links being up or down in each time-step. These steps are elaborated next.

1. **Generating Auto-Correlated Time Series:** The first step is to generate W independent time-series of link variables with desired auto-correlation. Different methods could be used for generating the time-series based on how the auto-correlation is specfied. We try two methods based on analysis of the field data we collected. The **first** method uses the long-range dependent characterisation of link connectivity we presented in Fig. 12. In this approach the link auto-correlation can be specified very succinctly by a single number: the Hurst parameter H, which for our experimental data was $H \approx 0.75$ across all links. We generate long traces of normalised fractional Gaussian noise (fGn) (with zero mean and unit variance) for this H, using the filtering method developed in [34]. The W fGn time-series thus generated each have the requisite auto-correlation properties; subsequent steps will cross-correlated them, and shift/scale them to have the appropriate link-specific mean and variance.
   The **second** method we use to generate the auto-correlated time-series assumes a linear stationary auto-regressive (AR) model [29] of appropriate order. An order p AR process derives the sample $x(t)$ at time-step t as:

$$x(t) = \sum_{k=1}^{p} a_k x(t-k) + w(t) \tag{5}$$

   The auto-correlation in the above process stems from the fact that the sample at time-step t is a weighted sum of the previous p samples, with an additional random noise component that has zero mean and constant variance. Based on the auto-correlation properties of links at different lags, we estimated that an AR process of order $p = 20$ matched our experimental data well. We then used the Yule-Walker method (`aryule` in Matlab) to estimate the AR coefficients, which were then applied as a filter to sequences of random white Gaussian noise to yield the desired auto-correlated time-series.
2. **Cross-Correlating the Time Series:** Having generated W sequences of independent variables with appropriate auto-correlations, this step introduces the cross-correlations as per the specified covariance matrix $\mathbf{C}$. The general idea is to take appropriate linear combinations of the W independent random variables to generate a new set of W random variables that have the desired cross-correlations. To this end we first determine the Cholesky decomposition of the covariance matrix, i.e. find the lower-triangular matrix $\mathbf{L}$ such that $\mathbf{C} = \mathbf{L}\mathbf{L}^{\mathbf{T}}$ where $\mathbf{L}^{\mathbf{T}}$ denotes the transpose of $\mathbf{L}$. The symmetric positive definite nature of $\mathbf{C}$ ensures that such decomposition exists and can be computed relatively easily (using `chol` in Matlab). However, for computation stability it is desirable to have $\mathbf{L}$ as sparse as possible. To this end we tried several methods to permute the rows and columns of $\mathbf{C}$ to make it more diagonally dominant, and

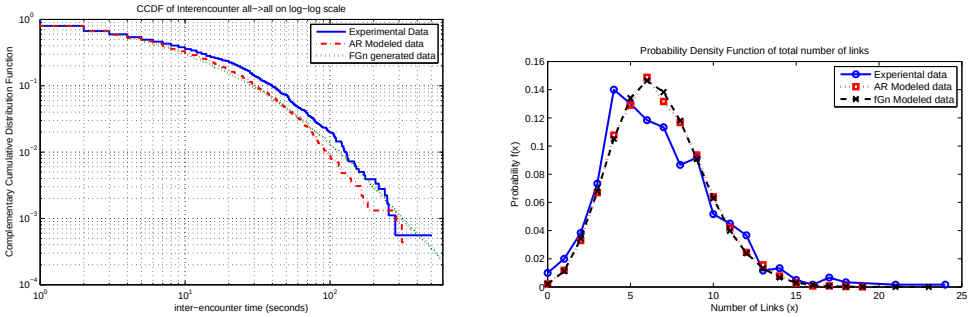

(a) CCDF of inter-encounter times

(b) PDF of total links

Figure 16: Comparison of model output with experimental data

chose the symmetric approximate minimum degree permutation (`symamd` in Matlab) to obtain the most sparse Cholesky decomposition $\mathbf{L}$.
Given a vector of W uncorrelated random variables $\mathbf{x(t)} = (x_1(t) \ \ldots \ x_W(t))$, we generate a vector of W correlated random variables $\mathbf{y(t)} = (y_1(t) \ \ldots \ y_W(t))$ with covariance as per matrix $\mathbf{C}$ using:

$$\mathbf{y(t)^T} = \mathbf{L x(t)^T} \tag{6}$$

where $\mathbf{L}$ is the Cholesky decomposition of $\mathbf{C}$.

3. **Converting to Binary Variables:** Random variable $y_i(t)$ above corresponds to link-i at time-step t, and already has requisite correlation with $y_i(t')$ (i.e. auto-correlation) as well as with $y_j(t)$ (i.e. cross-correlation with other links). In this step we convert continuous-valued $y_i(t)$ to corresponding binary values $z_i(t)$ by comparing with threshold T_i, i.e. $z_i(t) = 1$ if $y_i(t) > T_i$, and 0 otherwise. The threshold T_i is chosen so that

$$P[y_i(t) > T_i] = P[z_i(t) = 1] \tag{7}$$

The right side is the mean value $E[z_i(t)]$ of the link, which is available as input to the model. Random variable $y_i(t)$ is a linear combination of Gaussian variables $x_j(t)$, with weights known from the Cholesky decomposition matrix and the fGn/AR parameters, and therefore $y_i(t)$ is also Gaussian with known variance. Using tabulated values of the CDF of the normal distribution, the threshold T_i in Eq. (7) can easily be computed, and this threshold is then used for converting the model output to binary.

4.2.3. Validating the Model

To validate the model we compared its synthetic output with the empirical traces obtained from our experiments. Parameters estimated from the empirical trace from Game 1, such as mean and variance of each link, their auto-correlations, and the covariance matrix, were fed as input to the model. The

trace output by the model (i.e. binary time-series for each link) was subjected to the Kolmogorov-Smirnov (K-S) "goodness of fit" test and found to match the distribution seen in experiment for all links. Moreover, the statistical metrics directly controlled by the model, such as link mean, variance, auto-correlations, and cross-correlations, were found to match well.

We show that metrics that are by-products of the model (i.e. not directly part of the input) corroborate well with experiment. One such important metric is the inter-encounter time, which is a by-product of link auto-correlation. In Fig. 16(a) we show that the CCDF of the inter-encounter time seen in model trace data matches very well with the experimental dara from Game 1, confirming that our model has captured auto-correlations correctly. Another important metric is the number of links in the topology at any time-instant, which in turn is influenced by the cross-correlations amongst links. The PDF of the number of links shown in Fig. 16(b) again shows that our model matches well with experiment, affirming that the cross-correlations are also captured correctly by our model. Several other metrics (such as node degrees) seen in our model output were also found to match well with empirical data, and are omitted here for brevity.

4.2.4. Using the Model

Our model is fairly general: it takes as input the mean and variance of individual links, their auto- and cross-correlations, and outputs an arbitrary-length time-series of dynamic topologies with desired stochastic properties in terms of number of links, inter-encounter times, neighbour distributions, etc. Our model does not make any assumptions specific to the operating environment, and as such can be applied to model dynamic topologies arising in any mobile ad-hoc or delay tolerant network studies.

Deducing the input parameters to the model, in particular the cross-correlations between links, requires access to sufficient experimental data. Even then, estimating the parameters can be tricky: for example, the same soccer team plays each game differently depending on their strategy and their opponent. Nevertheless, we think reasonable approximations can often be made: for example, we can expect that links between players in similar positions (e.g. defenders) are more highly correlated with each other than with a link between players in different positions (e.g. defender and forward), or that a midfielder's link to a left-wing player will generally be negatively correlated with his link to a right-winger. We believe that capturing even a few key correlations (in a sparse covariance matrix) can give us much more realistic dynamic topologies for routing studies as compared to using overly simplistic models that ignore correlation effects altogether.

5. Multi-Hop Routing Algorithms for Real-Time Monitoring

Having understood the dynamics of the wireless topologies arising in the soccer field, in this section we evaluate routing mechanisms for extraction of

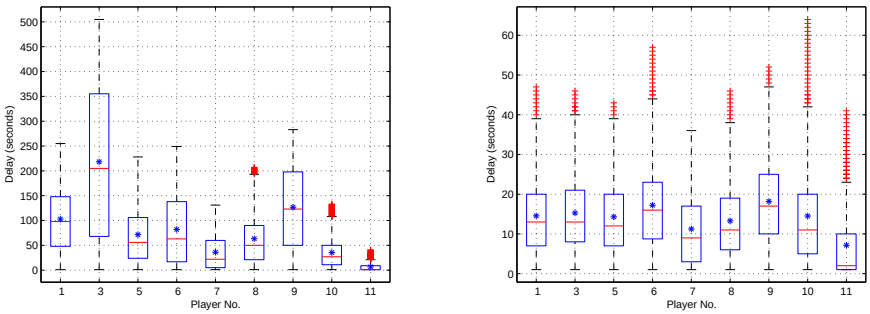

(a) Direct Delivery (b) Flooding

Figure 17: Comparison of delay values from Direct Delivery and Flooding (Game 1)

player data in real-time. In what follows we first argue that multi-hop routing is necessary for real-time performance, and then develop a practical yet efficient routing mechanism that achieves the desired trade-off between data delivery delay and energy performance.

5.1. The Need for Multi-Hop Routing

Using data from the monitored games, we first show that multi-hop routing has the potential to significantly reduce delays in extracting data from players in the field. We assume that the player-worn devices generate a sample (of physiological measures of the player such as heart-rate, ECG, oxygen saturation levels, impacts, etc) every second, and that every device in the field is given one transmission opportunity per-second (i.e. our MAC scheme is time division multiplexed with a periodicity of one second to avoid the possibility of collisions). We compare the 90-th percentile values of the delay in delivering a sample from an athlete-worn device to any of the base-stations under two schemes: (a) **Direct delivery**, whereby the player's transmission is received directly by a base-station, and (b) **Flooding**, whereby data from a player is forwarded by all recipients, and information propagates by one hop in each second till it reaches a base-station. Direct delivery has the lowest overheads (since it does not require store-and-forward routing), but is expected to incur higher delay, while flooding is delay-optimal (since all routing paths to base-stations are tried simultaneously) though resource-expensive (and hence impractical). Direct delivery and flooding therefore set upper and lower bounds on the achievable delay.

For Game 1, we depict in Fig. 17 a box-plot of the delay values of delivering data from each player to a base-station for direct delivery and flooding. In this experiment, we did not use high-gain antennas for the base-stations, so the receive range of the base-stations was small and consequently Fig. 17(a) shows that delays are large when we rely on direct delivery of data. It should

not be surprising that direct delivery delays heavily depend on playing position due to poorer radio connectivity at the centre of the field than the edges: for example midfielders (player 3) incur high delay (mean, depicted by a star, of about 220 seconds), defenders (e.g. player 10) lower delay (mean of about 40 seconds), while the goalkeeper (player 11) has close to zero delays. By contrast, Fig. 17(b) shows that with flooding, data delivery delays are reduced substantially, to within a mean of 20 seconds uniformly across all players. This clearly demonstrates that multi-hop routing has the potential to enable real-time monitoring of all players.

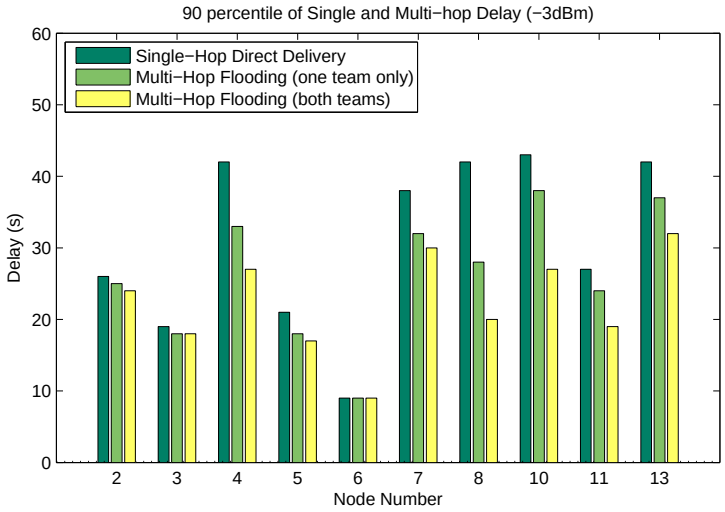

Figure 18: Comparing 90-th percentile delays from direct delivery, flooding using one team, and flooding using both teams (Game 2)

For Game 2, we recorded position information of all players (from both sides) throughout the game, using which we generated wireless connectivity using our model presented in section 3, as described earlier. This time we assumed that the base-stations have 12dBi-gain antennas (and hence extended reach), but we also assumed that the body-worn devices used transmit strength of -3dBm (rather than the 0dBm strength available on the MicaZ motes) in-line with the observation that true body-wearable platforms (such as the Sensium$^{\text{TM}}$[6]), being extremely small and lightweight, are limited in their transmit power. Once the second-by-second connectivity is generated (amongst players and to base-stations), we are able to evaluate data delivery delays as in the previous example.

In Fig. 18 we plot the 90-th perecentile values of the delays for each player obtained from direct delivery and from flooding (assuming first only one team, and then both teams, are outfitted with the monitoring devices). Some interesting observations can be made from the figure: (a) A midfielder (such as player

8) benefits immensely from multi-hop routing, in this case bringing delay from 42 down to 20 seconds (when both teams are outfitted), which is significant when considering a live broadcast of a soccer game to TV audiences. (b) The goalkeeper (node 6) and defenders (e.g. node 3) do not require multi-hopping of their data, since they are usually close enough to the edge of the field to have direct connectivity with the base-station (which has high-gain antenna). (c) A striker (e.g. node 10) benefits more from multi-hopping through the *other* team than through his own team, which should not be suprising since the striker can be expected to spend more time in the opponent's half. To summarise, we note that multi-hopping reduces delay significantly enough for some nodes (predominantly mid-fielders), and moreover makes delays more uniform across all nodes, which is important if monitoring of soccer players has to be done in real-time.

5.2. Routing Protocol Requirements

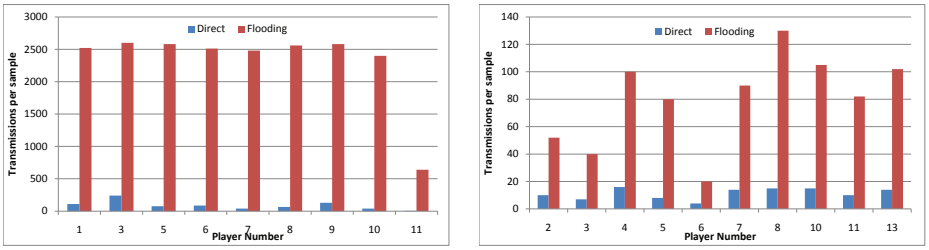

(a) Game 1 (b) Game 2

Figure 19: Mean number of transmissions per sample required by direct delivery and flooding for (a) Game 1 and (b) Game 2.

As noted earlier, direct delivery gives an upper limit while flooding gives a lower limit on delivery delays. Conversely, in terms of energy, direct delivery is much more efficient than flooding. This is because direct delivery requires fewer transmissions (bear in mind that multiple transmissions may still be required since the body device transmits data every second even if no base-station is within hearing range) compared to flooding, which requires transmission of many copies of each data item by many nodes. In Fig. 19 we plot the mean number of transmissions required for each sample of each player to be extracted using direct delivery and flooding. In Game 1 (small antenna base-stations), flooding requires nearly two orders of magnitude more transmissions per sample than direct delivery, while Game 2 (base-stations have high-gain antennas) shows about an order of magnitude difference in energy. This illustrates that flooding can reduce data extraction delays, but comes at the cost of increased energy requirements. In what follows, we develop a new multi-hop routing mechanism that allows the trade-off between energy and delay to be explcitly controlled.

In our prior work [35] we tried several single-copy and multi-copy routing schemes, but found them to have fixed tradeoffs points between delay performance and resource consumption. From an application viewpoint, we would

like to be able to *tune* the performance of the scheme, either towards better delay performance or lower resource consumption. To allow for good delay performance, we begin with a flooding-based approach. A practical flooding-based strategy must be able to limit both the memory requirements as well as the transmission bandwidth. The tunable scheme we propose allows us to control these parameters, and incorporates the following features.

5.2.1. Replication at the Source

In the proposed scheme, a player maintains a *window* of samples of length W for itself in a FIFO manner; thus, a newly generated sample pushes out the oldest sample in the window. During its slot, a player transmits its entire window of samples. Thus, every sample is *replicated* W times by the source. This guards against the case where a sample may be lost because no other player (or base-station) is able to hear the transmitter at that time.

5.2.2. Replication at Intermediate Nodes

Further, every player also maintains a window of W samples for *every other player*. In effect, this allows replication of the samples at *every intermediate hop*. When player i is not transmitting, it listens to other players' transmissions. Packets heard from other players are then analysed for data. If a received packet contains a more recent window of samples for player j, this data replaces the contents of the window for player j held by player i.

5.2.3. Data Freshness

The above two steps limit the amount of data transmitted in any packet to a $N * W$ matrix of samples (where N is the number of players). However, we still have the issue of old data being forwarded within the network. This may arise, for instance, when player i forwards its data to player j, and then becomes disconnected from other players for an ensuing interval. In this case, the data from player i may become obsolete (from an application point of view), yet continue to be forwarded by player j and its neighbours.

To ensure *data freshness*, player j forwards data for player i only if the *most recent* sample in its window for player i is less than A samples old. Thus, each intermediate node filters the data it receives from surrounding players before forwarding it on. The quantity A (denoting "age") can be selected based on the requirements of the application. For instance, for television broadcast enhancement where players' heart rate is streamed in real time, samples which are more than 10 seconds old may become irrelevant, and A can be set to 10.

5.3. Operation of the Proposed Protocol

As opposed to the forwarding-based schemes, our flooding-based scheme will actively *drop* a window of samples as they age beyond A. As a result, certain samples may never get delivered to the base. Therefore, as compared to schemes so far, we need to evaluate an additional parameter, the *delivery ratio* of this scheme.

Further, note that W controls the amount of replication at the source, while A controls the amount of replication in the network. A smaller value of A means that a given window of samples is likely forwarded through a smaller number of hops, as it would hit the age threshold faster. Hence small values of A result in lower delivery ratios (for players who are not well connected with the base station) as the samples are dropped within the network before they hit the base.

Finally, note that the *maximum* time a sample can spend in the network is $W + A$. Note that we have chosen to implement our algorithm in this form (rather than take a more traditional approach in terms of hop counts) for two reasons:

1. Using a window of samples for each player allows us to naturally control another important parameter, the *memory requirements* of this scheme, to a buffer of at most $N \times W$ samples per player.
2. Given that every player maintains a buffer of W samples for every other player, it is artificial to impose a hop count on every sample in that buffer. Rather, it is simpler and more intuitive to impose a maximum age A on the buffer, and to discard the entire buffer if the data is too old.

Finally, as with other schemes we assume that the base is able to inform a player when its transmission is received, upon which the player erases its entire window of samples.

5.4. Performance of our Protocol

We expect our tunable scheme to provide performance that can be varied between the extremes of direct delivery and full flooding. We therefore evaluate the delivery ratio, resource consumption and delay characteristics of our scheme, for different settings of parameters window-size W and age A.

For Game 1, we show in Fig. 20 the packet delivery ratio, mean delay, and resource consumption (i.e. mean transmissions per packet) for selected players (space constraints prevent us from including plots for all players) for various settings of the key parameters window-size W and age A of our algorithm. Our first observation is that the Goalkeeper (and in general players who are at the edges of the field and hence well-connected to the base-stations) have little need for multi-hop routing, and thus their delivery ratio is insensitive to A, and depends only on W. Similarly, their delay performance is insensitive to A since they deliver most of their packets directly to the base. Further, their resource consumption saturates with increasing W, since the buffer gets flushed when the data reaches the base, and samples are typically not transmitted for the entire duration of the window. By contrast, the figure shows that player 1 (Centre Midfielder) benefits from multi-hop routing: for small A (e.g. 5), the burden of deliver falls on the source, and the delivery ratio is low. As A increases, mean delay decreases and delivery ratio increases, showing that forwarding load is balanced between the source and the network. The resource consumption (mean packet transmissions per sample) increases roughly linearly with W, with a slope determined by A, which determines the degree to which a sample sent by the source is forwarded in the network. These results show that

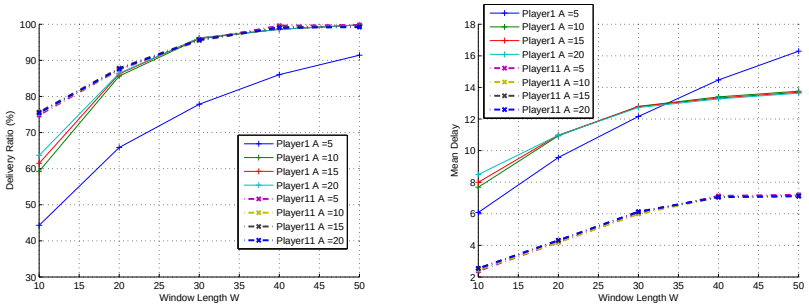

(a) Delivery Ratio

(b) Delay

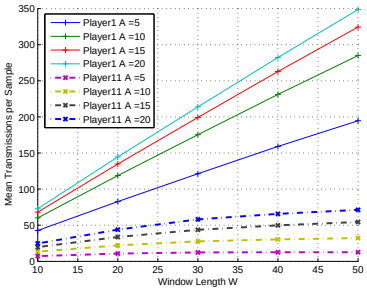

(c) Resource Consumption

Figure 20: Performance for player 1 (Center Attack) and player 11 (Goalkeeper) in Game 1 as a function of W and A: (a) Delivery ratio, (b) Mean delay, and (c) Energy consumption

by adjusting the values of A and W, it is possible to achieve desired trade-off between energy and delay performance.

Fig. 21 shows the performance of our routing scheme for various settings of parameters W and A for Game 2. Due to space constraints, we only depict results for Player 8 (Midfielder), who benefits well from multi-hop routing. Fig. 21(a) shows that the delivery ratio increases monotonically with A, for a given window size W, as the burden of delivery moves from the source to the network. The delivery ratio also increases monotonically with the window size W (for fixed A), since increasing replication by the source increases the likelihood of the packet reaching the base, regardless of the forwarding contributions from the network.

Fig. 21(b) shows the 90-th percentile values of the delays from our routing scheme (only packets that are delivered are consider for this computation). We note that if we set both W and A to be large enough ($W = A = 35$), the delivery ratio is very high (over 98%), while delays are about 20 seconds (which is identical to the delay of 20 seconds obtained by flooding as shown in Fig. 18), while resource consumption is about half that required by flooding (72 transmissions-per-sample versus 133 transmissions-per-sample), as shown in Fig. 21(c). Not only is our scheme much more efficient, but it also offers the flexibility to control

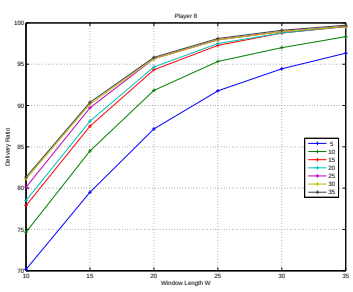

(a) Delivery Ratio

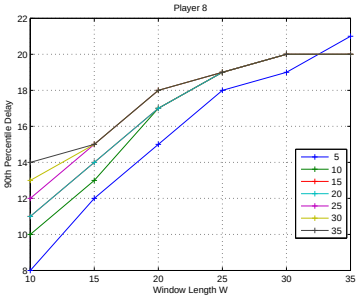

(b) Delay

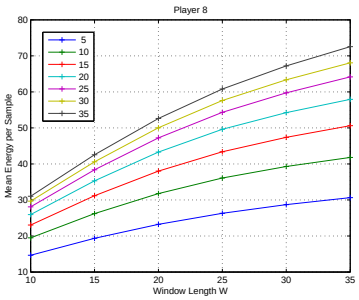

(c) Resource Consumption

Figure 21: Performance for Player 8 (Midfielder) in Game 2 as a function of W and A: (a) Delivery ratio, (b) Mean delay, and (c) Energy consumption

the performance trade-offs. It's resource consumption can be further restricted, for example by reducing the values of W and A (say $W = 10$, $A = 5$); the trade-off is that the delivery ratio will reduce to about 70%. Lastly, we note that there are multiple combinations of A and W that have similar resource utilisation but give different performance: for example for player 8, we can achieve the required delivery ratio of 95% by increasing source replication W ($W = 35$, $A = 5$) or by increasing network involvement ($W = 25$, $A = 10$). Using a higher A value increases the resource consumption faster than using a longer window (in this case, 30.5 transmissions-per-sample versus 36 transmissions-per-sample), since it increases the flooding within the network. However, it may not always be possible to achieve high delivery ratios using one parameter alone.

6. Conclusions and Future Work

Wireless sensor networks offer unprecedented ability to monitor athletes during competitive sporting events, enabling new applications for injury reduction, referee-assist, and enhanced TV broadcast services. This paper presents the first study of the opportunities and challenges arising in monitoring soccer players

during live games. Using experimentation with real soccer teams playing competitive games, we have developed profiles of wireless connectivity in the soccer field, characterising aspects such as neighbourhood, inter-contacts, and correlations. Using these profiles we have shown that current and emerging body-wearable platforms will not have adequate range for direct extraction of athlete data in real-time, and multi-hop routing will be required. We develop a novel yet practical routing scheme that allows the delay-energy performance trade-off to be tuned between the two extremes of direct transmission and flooding. Our work sets the foundation for future mobile sensor network systems for real-time monitoring of athletes in field sports.

Several directions for future work can be envisaged. We could always benefit from more data on mobility of players on the soccer field, including the position of the ball. More extensive data would help us validate robustness of key parameters of our connectivity model (such as link auto- and cross-correlations) across multiple games. We would also like to undertake experimental work that replaces Crossbow motes with emerging truly wearable body-monitoring devices (such as from Toumaz Technologies, VxSport, or RealTrack Systems) as they become available. The routing protocols can then be implemented on such body-worn devices to profile their real-time performance in live games.

References

[1] Armstrong, S. (2007) Wireless Connectivity for Health and Sports Monitoring: A Review. *British Journal of Sports Medicine*, **41**, 285–289.

[2] GP-Sports. SPI Elite. `http://gpsports.com`

[3] Sensixa-Ltd. e-AR. `http://www.sensixa.com`

[4] VxSport. VxLog Model 100 and 220. `http://www.vxsport.com`

[5] RealTrack Systems. WiMu: Physical activity monitor. `http://www.realtracksystems.com`

[6] Toumaz-Technologies-Ltd. Digital Plaster using Sensium^{TM}. `http://www.toumaz.com`.

[7] Power-Paper. Power Patch Platform. `http://www.powerpaper.com`

[8] Smith, R. and Loschner, C. (2002) Biomechanics Feedback for Rowing. *Journal of Sports Sciences*, **20**, 783–791.

[9] Hallberg, J., Svensson, S., Ostmark, A., et al. (2004) Enriched media-experience of Sport Events. *Proceedings of the 6th Workshop on Mobile Computing Systems and Applications*, December.

[10] Konberg, T., Ohult, C., and Delsing, J. (2003) Measuring breathing and heart rate data with distribution over wireless IP networks. *Instrumentation and Measurement Technology Conference*.

[11] Afonso, J., Silva, H., Oliveira, P., Correia, J., and Rocha, L. (2007) Design and Implementation of a Real-Time Wireless Sensor Network. *SENSORCOMM*, October, pp. 496–501.

[12] Sivaraman, V. Athlete Monitoring Project. `http://www2.ee.unsw.edu.au/~vijay/athlete`

[13] Crossbow Technology, Inc (2007) *MPR/MIB Users Manual*. `http://www.xbow.com`

[14] Schnayder, V. et al. (2005) Sensor Networks for Medical Care. Technical Report TR-08-05. Division of Engineering and Applied Science, Harvard University.

[15] Xiao, S., A.Dhamdhere, Sivaraman, V., and A.Burdett (2009) Transmission Power Control in Body Area Sensor Networks for Healthcare Monitoring. *IEEE Journal on Selected Areas in Communications*, **27**, 37–48.

[16] QStarz. BT-Q1300ST GPS Sports Recorder. `http://www.qstarz.com/Products/GPS%20Products/BT-Q1300ST-F.htm`

[17] Guterman, J., Moreira, A., Peixeiro, C., and Rahmat-Samii, Y. (2008) Comparison study of electromagnetic human interaction with various 2.4 ghz laptop integrated antennas, . may, pp. 1 –4.

[18] Hall, P. S. and Hao, Y. (2006) *Antennas and Propagation for Body-Centric Wireless Communications*. Artech House.

[19] Hao, Y., Alomainy, A., Hall, P., Nechayev, Y., Parini, C., and Constantinou, C. (2005) Antennas and propagation for body centric wireless communications. *IEEE/ACES International Conference on Wireless Communications and Applied Computational Electromagnetics*, april, pp. 586 – 589.

[20] Kurusingal, A., Dhamdhere, A., and Sivaraman, V. (2010) Modeling Signal Strength of Body-Worn Devices. *IEEE LCN*, Oct, . Denver, CO, USA.

[21] Crossbow technology - mica2 and micaz motes. `http://www.xbow.com`

[22] TP-link TL-ANT2412D Antenna. `http://www.tp-link.com/products/print.asp?pmodel=TL-ANT2412D`

[23] Chipcon. CC2420: 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver. `http://www.chipcon.com`

[24] Sivaraman, V., Grover, S., Kurusingal, A., Dhamdhere, A., and Burdett, A. (2010) Experimental Study of Mobility in the Soccer Field with Application to Real-Time Athlete Monitoring. *IEEE WiMob*, Oct, . Niagara Falls, Canada.

[25] Rhee, I., Shin, M., Hong, S., Lee, K., and Chong, S. (2008) On the Levy-Walk Nature of Human Mobility. *IEEE INFOCOM*, Apr, . Phoenix, AZ, USA.

[26] Cai, H. and Eun, D. (2008) Toward Stochastic Anatomy of Inter-Meeting Time Distribution under General Mobility Models. *ACM MobiHoc*, . Hong Kong.

[27] Chaintreau, A., Hui, P., Crowcroft, J., Diot, C., Gass, R., and Scott, J. (2006) Impact of Human Mobility on the Design of Opportunistic Forwarding Algorithms. *IEEE Infocom*, Apr, . Barcelona, Spain.

[28] Karagiannis, T., Boudec, J.-Y. L., and Vojnovic, M. (2007) Power Law and Exponential Decay of Inter-Contact Times Between Mobile Devices. *ACM MobiCom*, Sep, . Montreal, Canada.

[29] Box, G., Jenkins, G., and Reinsel, G. (2008) *Time Series Analysis: Forecasting and Control*, 4th edition. Wiley.

[30] Dallal, G. (2007). Historical background to the origins of p-values and the choice of 0.05 as the cut-off for significance.

[31] Bai, F. and Helmy, A. (2004) Wireless Adhoc and Sensor Networks. Kluwer Academic Publishers.

[32] Camp, T., Boleng, J., and Davies, V. (2002) A Survey of Mobility Models for Ad Hoc Network Research. *Wireless Communication and Mobile Computing (WCMC): Special Issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, **2**, 483–502.

[33] Tan, F., Borghol, Y., and Ardon, S. (2008) EMO: A Statistical Encounter-Based Mobility Model for Simulating Delay Tolerant Networks. *IEEE WoWMoM*, Jun, . Newport Beach, CA, USA.

[34] Ostry, D. (2006) Synthesis of Accurate Fractional Gaussian Noise by Filtering. *IEEE Trans. Information Theory*, **52**, 1609–1623.

[35] Dhamdhere, A., Chen, H., Kurusingal, A., Sivaraman, V., and Burdett, A. (2010) Experiments with Wireless Sensor Networks for Real-Time Athlete Monitoring. *IEEE SensApp*, Oct, . Denver, CO, USA.