The Smart Stadium Testbed for Integrating IoT Systems and their Applications

The Spring 2023 Smart Stadium VIP Team:
Team Leaders: G. Boudreau, R. Kelly, J. Paden, P. Raj, S. Sharma
Team Members: Y. Arora, L. Barros, A. Chen, H. Chen, N. Eliacin, A. Gao, E. Garber, A. Goli, A. Ismail, M. Madhusudhan, A. Marquez, A. Mehta, K. Metheny, S. Mittal, F. ”Finn” Morrison, K. Mutimer, S. Narayan, M. Patel, S. Patel, S. Punwani, R. Rajesh, B. Rudolph, J. Ryzhkov, N. Sethi, J. Sherling, A. Singh, P. Sreedhar, R. Sudunagunta, M. Um, N. Williams, W. Williams
Graduate Student: G. Lanier
Advisers: D. Phanish, P. Pinto, E. J. Coyle
VIP Program, Georgia Institute of Technology, Atlanta, Georgia 30332, USA
{ejc, dnagendra3}@gatech.edu, https://vip.gatech.edu/teams/vp3

***Abstract*—On a game day, sports stadiums are filled with: 50,000+ fans; medical, security, media, and venue personnel; and, the teams' coaches, players, and staff. It is thus an opportunity to research, design, test, and integrate many innovative sensing and communication systems. The Smart Stadium VIP team at Georgia Tech focuses on these opportunities during football games in Bobby Dodd Stadium. Our fundamental work in wireless sensor networks, cognitive radio networks, machine learning for play classification, and fan infotainment results in IoT and other systems that are designed and deployed in the stadium. Most importantly, all of our work is integrated, as a large-scale system, to improve fans' experiences of the game. This paper reports on the status of our research and the testbed and describes our goals for the future.**

***Index Terms*—mobile apps, wireless sensor networks, peer-to-peer communications, machine learning, commercialization**

I. Introduction

Almost every fan and staff member at a football game has a smartphone that has cellular, WIFI, and Bluetooth connectivity. Their phones also have accelerometers, microphones, cameras, and other sensors. The media organizations have high-definition video cameras, two-way radios, wireless links with reporters on the sidelines, etc. Each team often has its own wireless system to enable coaches on the field to talk with team personnel elsewhere in the stadium.

A stadium nowadays has infrastructure to handle the short-term but very large demands for communication and sensing from all of the people and systems at the game. It will have a distributed antenna cellular system, hundreds of WiFi access points, and bluetooth systems for sensing, tracking and short-range communications. It *will soon* have sensors for structural monitoring, localization of RF sources in the stadium [1], and audio monitoring of the venue for safety and other purposes.

Thus, in just 1 square kilometer, large sporting events have wireless communication and sensing potential of greater complexity than found almost anywhere else. They are thus an ideal environment in which to study and integrate these systems to produce the *Smart Stadium* of the future.

Fig. 1. Logo created by teammate Ryan Rodriguez for the *Smart Stadium*© VIP team [2].

We have thus developed the Smart Stadium testbed at Georgia Tech's Bobby Dodd stadium. This paper describes the systems in this testbed and discusses future enhancements. This testbed is *unique* because its systems are integrated with each other to enable new applications and capabilities:

- An enhanced wireless sensor network to measure how the stadium responds to crowd activities: Section II. These measurements enable both structural monitoring and characterization of the level of excitement of the fans. They are shared with our other systems for fan infotainment [3] and games in our new FANPLAY app.
- Our previously developed system [4] for enabling fans' on-demand access to any of the annotated video clips of plays collected as the game progresses.
- FANPLAY: A new suite of games for fans for entertainment before, during and after the game: Section III. Some of these games utilize the measure of fan excitement from the wireless sensor network. This suite will next include peer-to-peer games based on the clustered communication architectures in [5] [6].
- A system, designed via machine learning and stochastic modeling of play sequences, to automate the matching of video clips of plays with their appropriate annotation for the infotainment application in [4]: Section IV.

Our long-term goal is the commercialization of this system. In Section V, we describe the activities, such as surveys of fans,

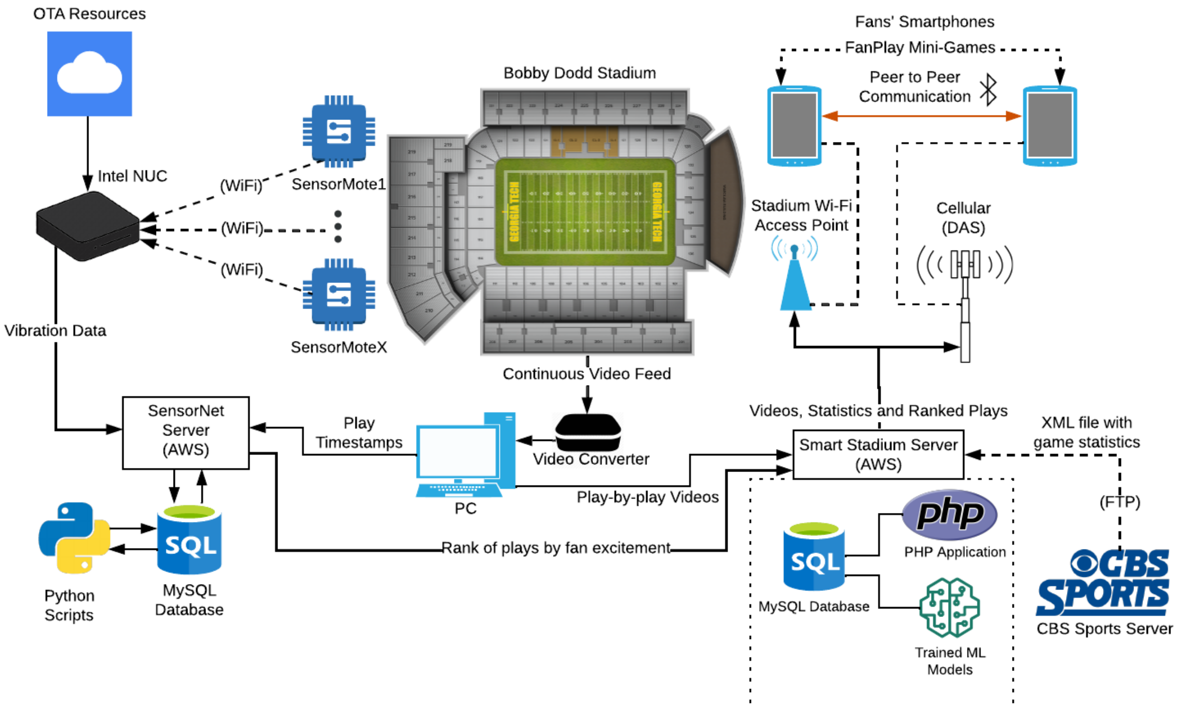

Fig. 2. Smart Stadium full system diagram.

that will identify strategies for monetizing the Smart Stadium system.

Success in these efforts is made possible by the Vertically Integrated Projects (VIP) Program [7]. The long-term, large-scale, multidisciplinary teams created by this program enable ambitious projects to be completed [8] [9].

II. Sensor Network

The SensorNet is a wireless IoT testbed with the goal of collecting data on fan interaction as well as monitoring a stadium's structural health. The system was last tested in 2016 and has been under a software and hardware design overhaul.

TABLE I
SensorNet Terms

Term	Definition
SensorMote	Wireless node deployed below stadium seating to capture sensor data
Cluster	A set of SensorMotes plus a base station to forward data to a server
Coordinator	Depricated, mictocontroller-driven decive used for hosting an access point for the previously-deployed SensorNet
ClusterHead	Base station hosting an access point and servers to interface with SensorMotes
Hypeness	Measure of fan excitement

A. Microcontroller and Network Updates

This overhaul was originally incited by hardware restrictions of the deployed microcontrollers. As stated in [3], the system tested in 2016 was limited to collecting vibrations on 2 axes (excluding the z-axis) due to a lack of analog-to-digital converter (ADC) pins on the Texas Instruments MSP430F5438A [10].

The SensorNet development team decided to transition to a new set of microcontrollers: the TI MSP-EXP432P401R, equipped with 24 ADC compatible pins and compatability with the CC3120MOD network card. The new wireless communication protocol being used is now Wi-Fi, as opposed to the previously implemented Zigbee system [10]. This, along with the newly available 32-bit timers on the MSP432, allow for better sync across the system, leading to better excitement measuring and capability to use other sensors such as one for audio later in the future.

The accelerometer of choice for testing the new system is the EVAL-ADXL354BZ, chosen due to the ultra-low noise spectral density (22.5 ug/Hz) to capture relatively small vibrations from stadium movement. A signal conditioning amplifier is placed before the ADC pins to boost the output range of the ADXL from 1.8 V to 3.3V, the expected reference voltage set by the MSP432.

B. Quality of Life and Ease of Use Improvements

Development has begun on a website for users to schedule SensorNet data collection. This website, using PHP with Envoy, will take in requested settings and communicate with all ClusterHeads to configure and schedule a start time for data collection. A user will be able to enter settings for the Cluster as well as add new nodes to the system. Cluster specific parameters include cluster number, period of time to collect data, and sleep time. SensorMote node parameters include application type (ex. vibration or audio data), data sampling rate, and sample count per WiFi packet. The website will be deployed on AWS and act as both a private portal for team admins to configure sensors, as well as a publicly facing website to display our work, sensor statuses, and other information.

The previous system was fully manual in its update and management process, requiring a full redeployment, and this was a major pain point for maintenance of the system. As a new feature, the team has been exploring over-the-air (OTA) updates for the SensorMotes. This system would allow for remote firmware updates, removing all manual steps in managing the system with the exceptions of initial deployment and hardware updates. The OTA update process works by partitioning the onboard memory in two: one for the previous firmware, and one for SHA integrity-confirmed, extracted firmware downloaded over a WiFi network. The SensorMote will then attempt to reboot pointing at the secondary memory partition and, if successfully rebooted, mark the previous firmware for overwrite during the next OTA update. If the reboot fails, the SensorMote will fall back to the previously functional firmware and continue operations until the OTA update is reattempted at a later point. The OTA update process is based on an example provided by TI for the MSP432s. Currently, the pre-WiFi component is complete with automation for the compilation of device firmware and packaging into the required .tar format. On event in the Git Repo, Git Actions have been configured to trigger the pre-WiFi component and notify ClusterHeads that an OTA update is ready for the SensorMotes. At the same time, the packaged OTA update will be transferred onto the ClusterHead to allow for the SensorMotes to pull the update down across the existing access point. However, there are blocking issues related to the SensorMote's ability to correctly pull the packaged update. Solutions to this are currently in development, but as a result, the OTA update process is not yet fully functional. OTA update theory of operation can be seen in Fig. 3.

Data collection servers have been transitioned from a local MySQL server onsite to using AWS, storing both raw accelerometer data and calculated hypeness data [10] [3]. By leveraging cloud databases, the data collected by the SensorNet can be widely accessible and removes the requirement for onsite or self-hosted servers to manage the data load. The goal of connecting the website to these databases is for future features including a graphical representation of SensorMote-collected data on the website and using hypeness data to drive some of the FanPlay application's minigames.

III. FanPlay Application

The Smart Stadium team has deployed several applications in the past that relate to the fan experience. One previous project was a web application that fetched data from a service from CBS Sports that provided real time game data in the form of text descriptions of plays and game statistics [3].

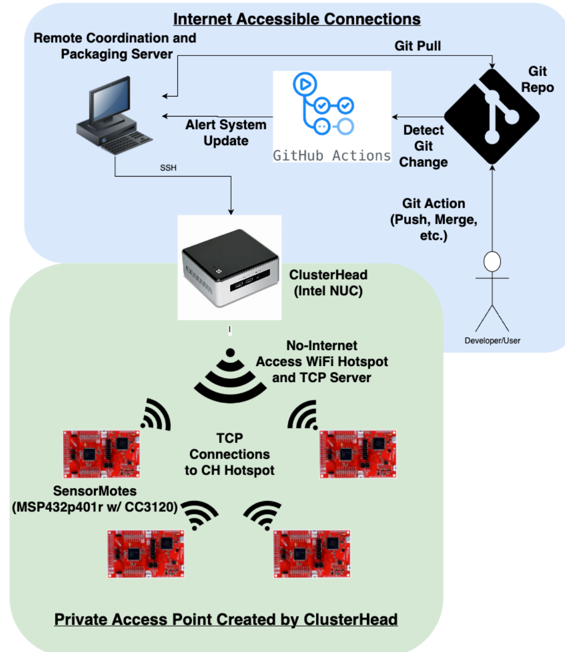

Fig. 3. SensorNet OTA updates procedure.

Members of the team would then manually link descriptions of plays with play recordings. The Machine Learning project is a continuation of this work.

Our current focus is FanPlay, a mobile application developed by Smart Stadium for use in athletic stadia during events. Our application currently includes several games for fans to play during athletic events as well as leaderboards. In the future, we also plan to have social features that allow fans in the stadium to connect with each other.

Behind the entertainment, social interaction, and games, we wish to stress the wireless communication networks in the stadium, and FanPlay is how we plan to do this. There is no better way to stress the network than by monitoring the effects of thousands of phones connecting to the networks. We also wish to collect information about the sensor network in the stadium using FanPlay.

During football games, thousands of people connect to a wireless network inside of the stadium at once. In our stadium, the sensors that we have deployed are also connected to this same network. This has the potential to stress the network and cause issues. We wish to learn more about these issues and learn how WiFi networks in small and medium-sized stadiums can handle abnormally large loads. This would provide insight for colleges, universities, high schools, basketball arenas, and many other environments.

We have a selection of games available for fans. The Prediction Game allows fans to predict the outcomes of athletic events, and they get points if their prediction is correct. Field Goal Frenzy is a simple game where players use their fingers to "kick" field goals and earn points. FanIQ is a trivia game where players answer questions and earn points.

One game that utilizes the sensor network is Hypemeter.

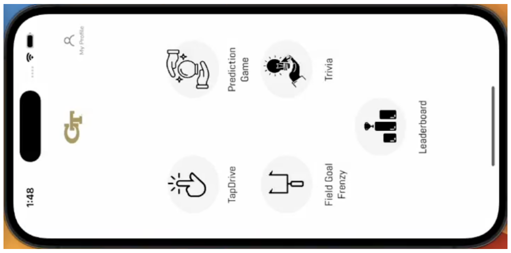

Fig. 4. FanPlay main menu for minigame selection. (shown horizontally)

This game allows all fans in the stadium to "compete" against each other. In essence, we are asking people in the stadium to get as loud and rowdy as they can, and our SensorNet will collect this information. We condense the information from the sensors into a "hype" value that can be narrowed down to specific sections of the stadium. The hype value is generated with the hypeness algorithm in [3].

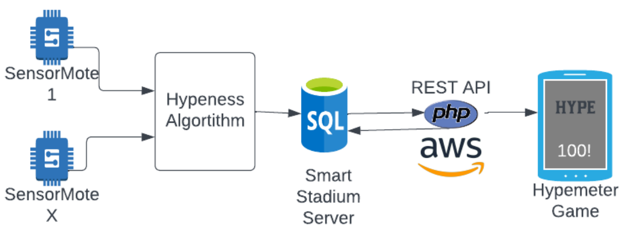

Fig. 5. Hypemeter game structure.

Another game that would use the SensorNet is Stadium Symphony. This game uses a configuration similar to Hypemeter, relying on fans to generate hype in the stadium. The stadium would participate in a musical game similar to Guitar Hero where individual sections generate hype at the correct times to hit musical notes and score points. The game would be displayed on fans' smartphones and the jumbotron. In addition to this, we have configured our server to accept accelerometer readings from individual mobile phones. We have not yet decided what we are going to do with this data, but it is an interesting data set to study.

We are also experimenting with the facilitation of multiplayer games, comparing utilization of external cloud servers to connections via Bluetooth or another wireless peer-to-peer (P2P) protocol. The P2P approach is promising because it gives us a lot more insight into how the stress on the wireless communication network in the stadium is impacted by the number of people connected to each other, especially since Bluetooth can interfere with 2.4GHz WiFi. This should presumably reduce the load on the stadium's wireless network during games, and free up capacity for other uses.

Another goal of the FanPlay app is to facilitate a means of social interaction amongst event attendees. We wish to allow people to connect to each other and exchange information while they are using our app in the stadium. Attendees will be able to share information like major, hometown, graduation year, greek affiliation, and other characteristics.

IV. Machine Learning

Complementing the fan engagement accomplished by the FanPlay application, the advanced analytics projects bridge the gap between gameplay and reported statistics along with developing innovative scouting strategies for the team. One such problem we are studying is: Can play classifications be done accurately using machine learning? Today, annotating each football play required several man-hours. If play annotations can be automated with minimal human intervention, it will help reduce the cost and increase user experience where we have labor shortage.

Our work incorporates elements of play videos and previous play data to create a deep learning model that combines the neural network feature vectors derived from play clips with a Markov Chain model of play sequences to classify the current play type: run, pass, field goal, etc. Results on Smart Stadium's machine learning work has been submitted for publication [11]. This topic is found to be of interest to the sports related machine learning research community [12].

The Smart Stadium team has procured a large dataset that includes high-definition videos of 5,647 Georgia Tech home-game plays manually recorded by the team over a decade, and statistical data from 56,496 play annotations in XML files created by the NCAA. The play annotations serve as ground truth for training our model. We have used GPU and high-memory nodes on Georgia Tech's PACE High Performance Computing center [13] to train and test our model. The model is based on the deep-learning shown in Fig. 6. It leverages transfer learning from well-established computer vision Recurrent Neural Network (RNN) with Long Short-Term Memory (LSTM) networks for extracting the feature vectors out of the video clips. For bench-marking, this standard base model, when trained/tested on the entire dataset, is able to classify plays with accuracies of [82.84, 88.56, 69.9, 70.83, 52.48] for [Kickoff, Run, Pass, ExtraPoint/FieldGoal, Punt]. The overall accuracy being 79.65%.

Major contributions in our work include the following items, further explained below:

- Understanding the outliers in play-by-play videos and labeling them for more accurate training.
- Creating a Markov Chain (MC) model for play prediction using the play annotations from previous plays.
- Exploring methodologies to combine the neural network and the Markov chain knowledge network to create the best Fusion model.

Data Clean-up: Identifying and cleaning defective data points is critical to machine learning. In analysing our play videos, we have identified 15 different error types, including video recording errors, very rare plays (i.e., an interception returned for a touchdown), plays obscured by ads, plays that started early or late in the video clip, etc. While some video clips can be fixed where the recording started too early, certain other where no play occurred or the recording started during

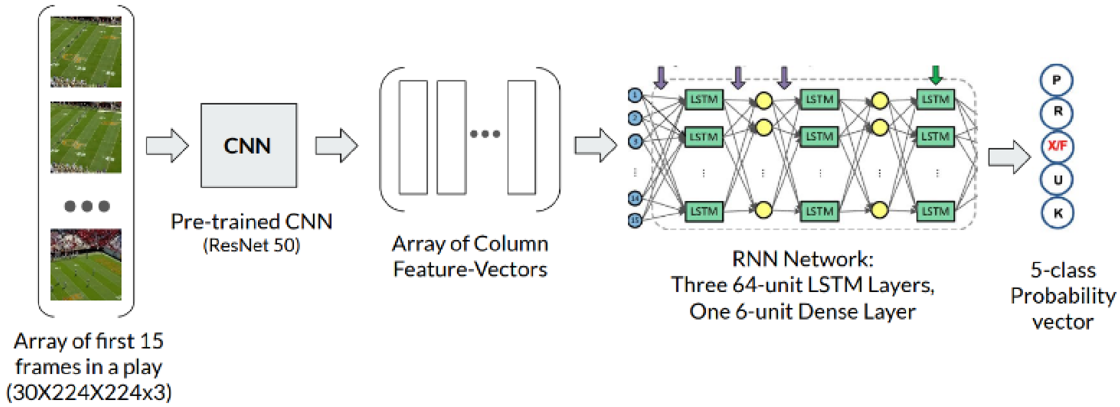

Fig. 6. Computer Vision Neural Network Model Flow Chart; [P = Pass, R = Run, X/F = Extra Point/Field Goal, U = Punt, K = Kickoff]. We consider a frame of the play clip for each of the first 15 seconds; Run the frames on the pre-trained ResNet50 architecture: 50 layer optimized convolutional neural network; Obtain output of feature vectors; Pass in feature vectors to the Long Short Term Memory Neural Network which outputs our final 5-class probability vector.

the play cannot be fixed and need to be purged from the dataset. After removing the defective clips, which reduced the dataset from 5,647 plays to 4,722 plays, the performance of the model improved by 2% to [85.96, 86.85, 70.33, 80.00, 59.79] with a total accuracy of 81.8%.

MC model: The 56,496 plays from the NCAA XML files are used to create a Markov Chain (MC) by extracting the state transition probabilities from real-world data. We have parsed XML files from 277 GT football games from 2000 to 2021 to extract statistics and play data for each play corresponding to the video clip. The MC state comprises of these 4 factors:

- Down (4): [1-4]
- Type (7): [Run, Pass, PAT, Field Goal, Punt, Kickoff, Other]
- Field Position (10): 100 yard field divided into 10 10-yd buckets on the home side: [0,9], [10,19] . . . [40,50] and opposition side: [50,40], [39,30] . . . [9,0]
- Distance to first down (6): [1,3], [4,6], [7,9], [10,14], [15,19], [20,99]

This gave us 4 * 7 * 10 * 6 = 1680 distinct possible states for a play in a game, yielding a 1680x1680 state transition matrix. Aggregate vectors are added for each row of the matrix that gives a probability vector of each play state leading to a [Run, Pass, PAT, Field Goal, Punt, Kickoff, Other]. This MC model is combined with the base Neural Network model to augment it with the game state information. The system is realistic since previous play sequences would be determined during the live game for which our system is running, and our goal is to classify the current play given previous plays.

Fusion model: We have explored four different ways of combining the NN results with the MC model. 1) We used thresholding and weighted sum of the prediction probabilities form the two models to produce a new probability vector. 2) We added the past play encoding as an input in addition to the 30x2048 video clip data to our LSTM network model. 3) We built a simple Sequential Network consisting of multiple Dense layers which takes in the result vectors of the LSTM model and the Markov Model. This approach faced a sample size issue since we could only use the probability outputs from the test set of the LSTM Neural Network when training the new network that combines the predictions. 4) We built a monolithic Neural Network, incorporating all of the play clips and corresponding MC state information spliced into the final dense layers. Table II compares the performance of these fusion methodologies. As seen, method 4 performs the best with individual accuracies of [89.29, 89.65, 83.42, 87.36, 83.10] and an overall average accuracy of 87.26%. This is an increase of nearly 7.6% from the original full-data base model accuracy. Each of these approaches are explained in further detail in the Stadium-IoPT machine learning paper [11].

TABLE II
TABLE COMPARING THE ACCURACY OF THE 4 FUSION MODELS

	Kickoff	Run	Pass	XP/FG	Punt	Total
MF 1	88.1	79.7	70.07	71.43	76.84	76.6
MF 2	82.14	88.31	76.33	85.71	69.47	83.1
MF 3	86.9	86.16	73.78	87.36	77.46	84.05
MF 4	89.29	89.65	83.42	87.36	83.1	87.26

The overall increases in accuracy with the above strategies is shown in Fig. 7. The significant increase in performance of the neural network when integrated with a Markov chain model shows that a knowledge network capturing the context of a play in addition to raw video data plays a vital role in the prediction. Our future work on this problem include: 1) Exploring other computer vision models beyond transfer learning for the deep neural network component. 2) Including a broader range of dataset using play videos from other colleges and NFL for achieving generalization. 3) Obtaining more play data that can enhance the accuracy of transition probabilities within the Markov Chain. 4) Exploring conflation techniques

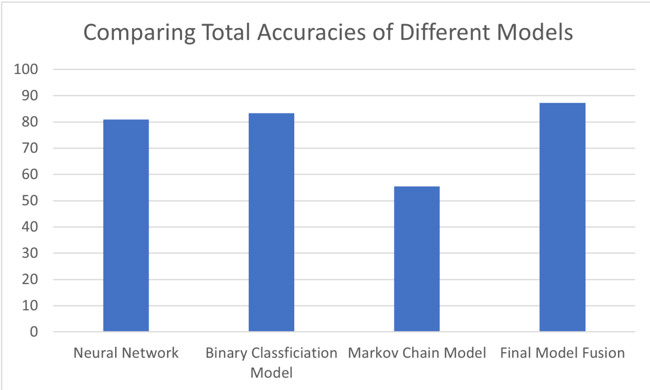

Fig. 7. Comparison of the best iteration of the Neural Network, Binary Classification Model, Markov Chain, and our best attempt at Model Fusion.

to combine the prediction probabilities from machine learning and knowledge network models.

V. Commercialization

A goal of our project is commercialization of the Smart Stadium system. We have looked into patents [14] for the FanPlay app and for the SensorNet or machine learning systems [15] to determine that we are not violating any existing patents. We are also planning to conduct demos of and surveys about the FanPlay app to gauge interest in it within Georgia Tech.

Fig. 8. Business Card with copyrighted Smart Stadium Logo and QR Code.

We believe the FanPlay apps will improve the game day experience for all attendees of Georgia Tech's athletic events. The demos and surveys will determine if this is true and help us gauge the best monetization method for the system. For example: (1) a willingness of fans see ads allows us to monetize FanPlay by sponsorships or inclusion of freemium content; (2) the FanPlay app could be offered as a free perk to incentivize purchase of Season Tickets.

We also plan to market the SensorNet technology to sports facilities, entertainment venues, fairs, and other places where crowds gather. We would license the rights to implement the technologies mentioned in this paper to organizations which would also bring in revenue through contracts. The possible valuations of the SensorNet are still in research by our team.

Conclusion

We described the current state of and future plans for the Smart Stadium testbed at Georgia Tech. The systems we developed show the many opportunities for innovation in an environment that is similar to what integrated communication and sensing systems will be like everywhere in the future.

Acknowledgment

We thank Georgia Tech Athletics for 10 years of access to Bobby Dodd Stadium for the deployment of our SensorNet and for gathering video clips and annotations of plays during home football games. We also thank Smart Stadium team members from prior semesters for their many contributions.

References

[1] S. Challa, et al, "Cognitive wireless networks: Research, applications, and testbed," in *Submitted to the 2023 IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2023.

[2] GT VIP Program. Smart stadium vip team description. [Online]. Available: vip.gatech.edu/teams/vp3

[3] H. Li, P. H. R. Pinto, D. Phanish, R. T. Abler, and E. J. Coyle, "Good vibrations: An iopt testbed for infotainment in a football stadium," in *2020 IEEE 6th World Forum on Internet of Things (WF-IoT)*, 2020, pp. 1–6.

[4] A. Ault, J. Krogmeier, S. Dunlop, and E. Coyle, "estadium: The mobile wireless football experience," in *2008 Third International Conference on Internet and Web Applications and Services*, 2008, pp. 644–649.

[5] D. Phanish and E. Coyle, "Application-based optimization of multi-level clustering in ad hoc and sensor networks," *IEEE Transactions on Wireless Communications*, vol. 16, no. 7, pp. 4460–4475, 2017.

[6] S. Laitrakun, D. Phanish, and E. Coyle, "Energy-efficient clustering and collision-aware distributed detection/estimation in random-access-based wsns," in *Data Fusion in Wireless Sensor Networks: A Statistical Signal Processing Perspective*, P. S. R. D. Ciuonzo, Ed. London UK: IET, 2019, pp. 129–158.

[7] E. Coyle, J. Krogmeier, R. Abler, A. Johnson, S. Marshall, and B. Gilchrist, "The vertically integrated projects (vip) program: Leveraging faculty research interests to transform undergraduate stem education," in *Transforming Institutions: Undergraduate STEM Education for the 21st Century*, G. Weaver, W. Burgess, A. Childress, and L. Slakey, Eds. West Lafayette, IN: Purdue University Press, 2016, pp. 223–234. [Online]. Available: http://www.jstor.org/stable/j.ctv2x00vcx

[8] B. Aazhang, et al, "Vertically integrated projects (vip) programs: Multidisciplinary projects with homes in any discipline," in *2017 ASEE Annual Conference & Exposition*, no. 10.18260/1-2–29103. Columbus, Ohio: ASEE Conferences, June 2017, https://peer.asee.org/29103.

[9] The VIP Consortium. A non-profit supporting vip sites around the world. [Online]. Available: vip-consortium.org

[10] D. Phanish, P. Garver, G. Matalkah, T. Landes, F. Shen, J. Dumond, R. Abler, D. Zhu, X. Dong, Y. Wang, and E. J. Coyle, "A wireless sensor network for monitoring the structural health of a football stadium," in *2015 IEEE 2nd World Forum on Internet of Things (WF-IoT)*, 2015, pp. 471–477.

[11] R. Pranav, et al, "Smart stadium: College football play classification," in *Submitted to the 2023 IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2023.

[12] S. Chen, Z. Feng, Q. Lu, B. Mahasseni, T. Fiez, A. Fern, and S. Todorovic, "Play type recognition in real-world football video," in *IEEE Winter Conference on Applications of Computer Vision*, 2014, pp. 652–659.

[13] PACE, *Partnership for an Advanced Computing Environment (PACE)*, 2017. [Online]. Available: http://www.pace.gatech.edu

[14] J. Kwon, "System for sport mornitoring based on internet-of-things," South Korea Patent KR1 020 200 118 166A, July 27, 2021.

[15] Q. Zeng, J. Liu, D. Yang, Y. He, X. Sun, R. Li, and F. Wang, "Machine learning based automatic sport event detection and counting," in *2021 7th IEEE International Conference on Network Intelligence and Digital Content (IC-NIDC)*, 2021, pp. 16–20.