Connected Cows: Utilizing Fog and Cloud Analytics Toward Data-Driven Decisions for Smart Dairy Farming

Mohit Taneja*†, Nikita Jalodia*†, Paul Malone*, John Byabazaire*†, Alan Davy*†, Cristian Olariu§
*Emerging Networks Laboratory, Telecommunications Software and Systems Group,
Department of Computing and Mathematics, School of Science and Computing,
Waterford Institute Of Technology, Waterford, Ireland
†CONNECT — Centre for Future Networks and Communications, Dublin, Ireland
§Innovation Exchange, IBM Ireland, Dublin
{mtaneja, njalodia, pmalone, adavy}@tssg.org, byabazairej@gmail.com, rioderaca@gmail.com

***Abstract*—The Internet of Things (IoT) is about connecting people, processes, data and things; and is changing the way we monitor and interact with things. An active incorporation of Information and Communication Technology (ICT) coupled with sophisticated data analytics approaches has the potential to transform some of the oldest industries in the world, including dairy farming. It presents a great opportunity for verticals such as the dairy industry to increase productivity by getting actionable insights to improve farming practices, thereby increasing efficiency and yield. Dairy farms have all the constraints of a modern business — they have a fixed production capacity, a herd to manage, expensive farm labour and other varied farm-related processes to take care of. In this technology-driven era farmers look for assistance from smart solutions to increase profitability and to help manage their farms well.**

We present an end-to-end IoT application system with fog assistance and cloud support that analyzes data generated from wearables on cow's feet to predict anomalies in animal behaviour that relate to illness such as lameness. The solution leverages behavioural analytics to generate early alerts towards the animals' well-being, thus assisting the farmer in livestock monitoring. This in turn also helps in increasing productivity and milk yield by identifying potential diseases early on. The project specializes in detecting lameness in dairy cattle at an early stage i.e. before the visual signs of it appears to be seen by farmer or an animal expert. Our trial results on a real world smart dairy farm setup consisting of a dairy herd of 150 cows in Ireland demonstrate that the designed system delivers a lameness detection alert up to 3 days in advance of manual observation.

***Index Terms*—Fog Computing, Cloud Computing, Internet of Things (IoT), Smart Dairy Farming, Data Analytics, Microservices, Data-Driven, Animal Welfare**

I. Introduction

The concept of smart dairy farming is no longer just a futuristic concept, and has begun to materialize as different fields such as machine learning have made inroads towards successful applications in this domain. The data-driven approach is transforming many industry sectors including dairy, and presents us with an opportunity to predict, control and prevent certain undesirable events.

The demand for dairy products is rapidly rising due to an ever increasing population coupled with an increase in income per capita [1]. Milk and dairy product consumption is higher in developed countries than in developing nations, but this gap is reducing with increasing incomes, rise in population, urbanization and dietary changes [2]. It has been estimated that the consumer base of dairy and dairy products is set to rise from 1.8 billion people in 2009 to 4.9 billion by 2030 [3]. However, methods to improve yield from the agricultural and dairy sector have not advanced at the same rate as the increase in demand. To cope with the increased demand for food, new and effective methods are required to increase the production capacity of this sector. Data-driven decisions, methods and measures can help in increasing the production capacity of these industries.

It can be expected that opting smart dairy farming principles which unify Internet of Things (IoT), data analytics, fog computing, and cloud computing will help meet these demands and contribute to sustainable growth in the dairy industry. The objective of the work presented is to enable data-driven decisions for dairy farming, and extracting timely insights from the data by designing suitable analytics models for such use-case scenarios. This aims to provide a set of controls to the farmer and other stakeholders to increase productivity, thus leading towards improved farming practices for the overall benefit of the industry.

The rest of the article has been organized as follows: §II presents the problem space being addressed, §III presents the real-world IoT smart dairy farm test-bed deployment, associated challenges, critical decisions, and experience gained throughout the process, §IV presents the design and development methodology used in building the end-to-end IoT solution, §V gives technical description of the solution with associate challenges and developed solutions, §VI presents the benefits to stakeholders, §VII presents the conclusion, §VIII presents ongoing and future work.

II. The Problem — Early Detection of Lameness in Dairy Cattle

Dairy farmers work hard from dawn till late in the evening — milking, feeding and maintaining the farm. So, it is a

Fig. 1: Cows with Long Range Pedometer (LRP) attached on front leg as part of our smart dairy farm setup.

challenge to monitor the well-being of hundreds of cows in the dairy farm in real time. The methods for looking after animal welfare are based on millennia of human experience and grounded on observational methods to analyse animal behaviour by visual observation for some kind of anomaly or potential health-issue. This leads to the question — *Could technology help? Why can't there be a better way to do it?*

There are behavioural changes when animals become ill, which can be mapped to specific illnesses. The risk of diseases has a large effect on the economy of a farm — payment for veterinary treatments and loss of milk production from the infected animals, as well as on animal welfare. *So, what if one could detect the onset of common diseases before any symptoms are even visible?*

To reiterate, the health and welfare of dairy cows is paramount to the productivity of the herd in both operational and capital expenditure related to pasture management and milk production. One of the issues that need to be addressed in this domain is lameness management.

Lameness is a condition that affects the locomotion patterns of livestock. An all-encompassing definition of lameness includes any abnormality which causes a cow to change the way that she walks; and can be caused by a range of foot and leg conditions triggered by disease, management or environmental factors. Controlling lameness is a crucial welfare issue, and is increasingly an inclusion in welfare assurance schemes.

Lameness is considered to be the third disease of economic importance in dairy cows after reduced fertility and mastitis [4]. It is estimated [5] that lameness costs an average of €275 in treatment per instance. Early lameness detection allows farmers to intervene earlier, leading to prevention of antibiotic administration and improvement in the milk yield as well as saving on veterinary treatment for their herd.

The existing solutions for lameness detection in dairy cattle either have high initial setup costs and complex equipment, or, in the ones that are technology based, major interoperability issues towards compatibility with existing farm based management solutions. As a solution to this, we have developed an end-to-end IoT application that leverages advanced machine learning and data analytics techniques to monitor the herd in real-time, and identify lame cattle at an early stage.

III. Real-World Testbed Deployment towards Smart Dairy Farm: Challenges, Decisions and Experience

Focused on animal welfare and health monitoring, this deployment involves installing sensors on cows' feet. Data generated from these sensors is subjected to analysis using fog computing, which is further enhanced by its cloud component that acts as the site for data fusion and other related resource demanding data analytics functionalities.

The trial was performed on a dairy farm having a herd size of 150 cows in Waterford, Ireland. The important decisions made during deployment and design phase of presented IoT solution have been listed in this section.

Decision 1: *Which wearable sensor technology to choose from the numerous options available for livestock monitoring?*

From the options available for the sensors/wearable for livestock monitoring, we decided to use radio communication based Long Range Pedometer (433MHz, ISM band) instead of WiFi based. The reason behind this was that the former does not depend on the Internet for its operation, and serve the purposes of data acquisition in farms where network connectivity is a constraint.

These wearables have less operational expense, and do not use WiFi based connectivity to send sensed data to a base station. Therefore, as a part of the real-world deployment, off-the-shelf available Long-Range Pedometers (LRP, ENGS Systems©®, Israel) specially designed for livestock monitoring were attached to the front leg of cows as shown in Fig. 1. A detailed analysis of other available options and previous approaches have been presented in [6].

The workflow and different components of the developed IoT solution have been presented in Fig. 2. These pedometers have sampling frequency of 8 milliseconds and forward their sensed data every 6 minutes. The sensed acceleration data is collected at a PC form factor device (fog node), where it is aggregated, pre-processed and converted into behavioural activities like step count. The system works in both housed and pasture based dairy systems. The cows are monitored continuously— whether they are in the fields during favorable weather conditions, or in-house during adverse weather conditions.

In this study, we used three behavioural activities (step count, lying time, swaps) for the analysis with their description as below:

1) *Step count:* This is the number of steps an animal makes.
2) *Lying time:* The number of hours an animal spends lying down, resting.
3) *Swaps:* This is the number of times an animal moves from lying down to standing up.

The choice of these three parameters is based on literature survey, which suggests that these three act as the best predictor of a lame cow or one transitioning to lameness while analyzing movement or activity patterns of cows.

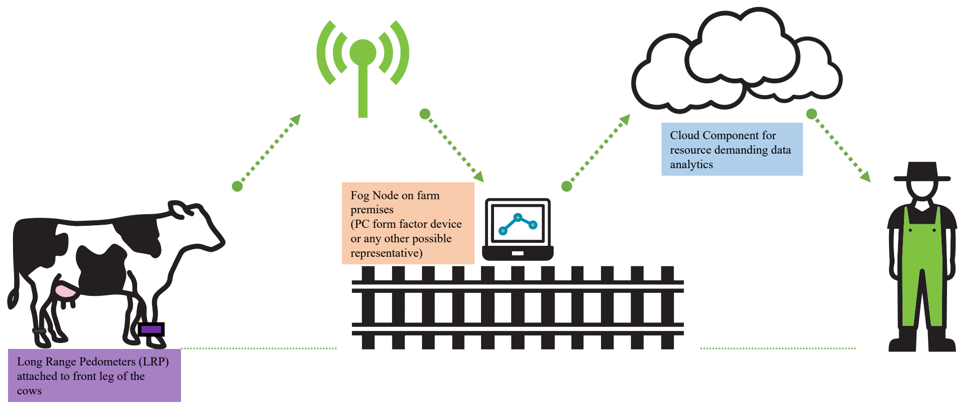

Fig. 2: System workflow, and different components of the developed IoT solution.

Decision 2: *Which network device amongst the available options along things to cloud continuum should be leveraged as a fog node in such IoT deployments?*

Fog computing is an emerging computation paradigm that aims to extend cloud computing services to the edge of the network, thus enabling computation closer to the source of data. It has been used increasingly in IoT applications, especially in constrained network and Internet connectivity scenarios, which is also one of the issues in remote farm-based deployment such as ours.

Most IoT enabled smart farms have some sort of farm management system in place which usually runs on a PC form factor device available within farm premises. Farmers use it to maintain logs and to keep other details electronically at hand. So our plan was to utilize the computing resources already available in such scenarios and leverage them under the fog computing paradigm. Thus, we choose the laptop available with farmer in our case as the fog node. It should be noted that the developed system is fully capable to adapt if the fog node is changed to any other possible representative such as a gateway device. A detailed discussion on this aspect of the system, and also on using resource constrained devices with low computational power as fog node has been presented in [6].

This decision also helps to improve fault tolerance, and build up the system resilience to variable farm environments such as weather-based network outages and connectivity issues because of geographical remote locations of farms. In scenarios with low/no internet connectivity it becomes ideal to process the data locally as much as possible and send the aggregated or partial outputs over the Internet to the cloud for further enhanced analytical results. The fog computing based approach leads to effective utilization of available limited bandwidth and reduces the dependency on the cloud by facilitating part of data analytics involved in the solution on the network edge. A detailed description of the distribution of services and computational processes running on the edge and in the cloud for the presented solution has been described in [6].

Decision 3: *Which streaming protocol to use for streaming data from the fog component to the cloud component?*

There are a number of options available when it comes to streaming the data, for e.g. MQTT (Message Queue Telemetry Transport), AMQP (Advanced Message Queuing Protocol), XMPP (Extensible Messaging and Presence Protocol), etc. Each of these have their individual pros and cons, and selecting one depends on the use-case, objective and IoT deployment scenario.

Our aim was to use a lightweight protocol that can work in our use-case and is also widely supported by both academia and industry in such scenarios. After evaluating and comparing the available options, we selected MQTT as the connectivity protocol in our deployment. It is a lightweight, open-source, publish-subscriber model based protocol working on top of the TCP/IP stack, originally invented and developed by IBM [7].

Decision 4: *What should be the development design of the system so that it should be usable, compatible and able to serve in both user possible scenarios listed below:*

1) *when a farmer acts as the end-user?*
2) *when an agri-tech service provider acts as the end user?*

The end user in our scenario could be a farmer with an existing system or an agri-tech service provider who wants to provide more services to their clients. With that in mind, we decided that the system should be developed as 'Application/Software as a Service'(AaaS/SaaS), which can be used by the service providers to integrate with their existing systems or can be used directly by farmer as well.

This brings us to our next questions — *which software development technique (or architectural style) to use while developing the system?* The answer and discussion on this has been presented in greater detail in the next section.

IV. Designing and Developing Software Systems in Fog enabled IoT Environments with Cloud Support

Decision 5: *Which software architecture or software development methodology to use so that the designed system should be multi-vendor inter-operable, and should also be in-line with the finalized design of AaaS/SaaS mentioned above?*

Designing and developing software systems is an intricate process which requires profound understanding of the procedure, consideration of software architecture and development techniques involved, and knowledge of various interconnected components in the deployed physical or virtual infrastructure.

Microservices architectural style comes as the first realization of a service-oriented architecture and is currently in wide use by industry for software development and deployment as part of best DevOps practices. Given its successful and wide adaptation in cloud computing domain, a microservices based architecture seems quite an obvious candidate for use in such fog enabled IoT deployments, but its use is not straightforward. The design and operational practice is sometimes quite different between these two technological paradigms [8]. The major reason for this can be that the microservices approach comes from a different perspective, which is to efficiently build and manage complex software systems, which in turn came to realization as a move towards architectural modularity. The main drivers of modularity are: 1. Agility 2. Testability 3. Deployability 4. Scalability 5. Availability

The challenge now is how to apply the microservices approach to build the application in an IoT scenario leveraging fog computing paradigm. In our analysis, we found that a distributed modular application architecture using microservices was the best approach, given we could align with the service-based and event-driven needs of our application. Modularity is a must, though not every portion of production has to be a microservice. Microservices need collaboration, and only when there are one or more drivers present should one make use of microservices. In the our use case scenario, we had all of the above drivers present. Microservices come with a set of advantages that make it an ideal architectural style for software development in end-to-end IoT solutions with constrained environments, giving the ability to overcome the constraints of vendor lock-in, while attributing technological independence between each set of services that make up an application.

Thus, with this understanding we decided on following a hybrid microservices based approach for application design and development in our end-to-end IoT solution. This decision was also made keeping a future vision in mind of the work, where the microservices act as facilitators to enable dynamic service migration based on the network characteristics to increase quality of service and for better service provisioning.

V. Technical Challenges and Solutions

Data Analytics and Machine Learning

This section presents details on challenges faced and solutions developed while designing a machine learning model for animal behaviour analysis for early lameness detection in dairy cattle.

1) **Cow Profiles:** *How to build robust cow profiles that are distinguishable by the learning model as lame and non-lame? Which parameter to use as baseline while building and comparing cow profiles?*

For the system to differentiate between normal and anomalous behaviour due to lameness, we must first form profiles to characterize normal (non-lame) and lame behaviour in the herd. The most frequently used approach for this is to examine the activity level of lame and non-lame animals and study how

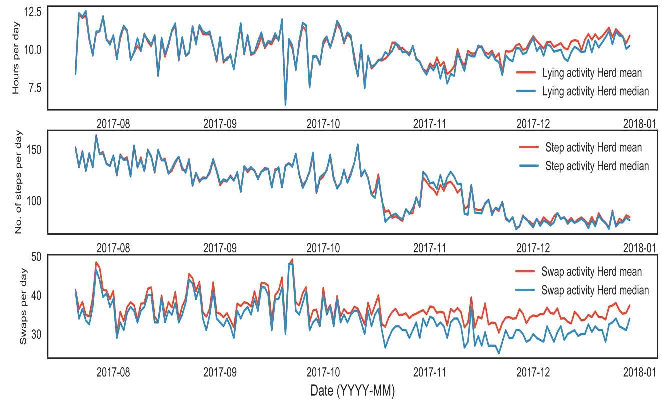

Fig. 3: Comparing the Mean and Median of the various Animal Activities

these differ from the mean of the entire herd. But as it is known that outliers (i.e., a single element in sample being too high or low) can affect the mean value of sample; hence median or quantiles are sometimes taken as a better measure. To address this issue, we studied the relationship between the herd mean and the herd median. The results of this as presented in Fig. 3 show that these almost trace out each other for all the three activities (Lying time, Step count and Swaps). This is one of the features of a normal distribution, and therefore it would not matter whether the mean or median is used. Thus, we decided to use herd mean in our analysis.

A study [9] on animal behaviour analysis and association patterns of cattle shows that animals grazing within the same pasture can influence the movement, grazing locations, and activities of other animals randomly, with attraction, or with avoidance; therefore most of the animals will have their activity levels almost equivalent to the herd mean.

For such reasons, using herd mean as the baseline seems appropriate. Thus, any deviation from the herd mean should serve as a preliminary indicator for a sign of change in behaviour, which could potentially be lameness, among other reasons. Such an analysis eliminates the effects of external factors, as these will be largely affecting the herd as a whole. Further, the measure used to note the deviation in behaviour while forming Lame and Normal profiles of cows in the herd was Mean Absolute Deviation (MAD), and while comparing behaviour of individual cows with these formed profiles was Average Deviation.

We build a profile for each animal to characterize normal behaviour in a time window using activity based threshold clustering, details of which have been presented further in the article. This helps us to define Lameness Activity Region (LAR, the period during which the animal is confirmed lame) and Normal Activity Region (NAR, the period during which animal is confirmed as non-lame), which later acts as ground truth input for the classification model for detecting lameness. An example of this has been presented in Fig. 4 for a random cow with ID 2346 in the herd.

However, by comparing the activity of each cow against the herd mean, we found out that not all animals behave the

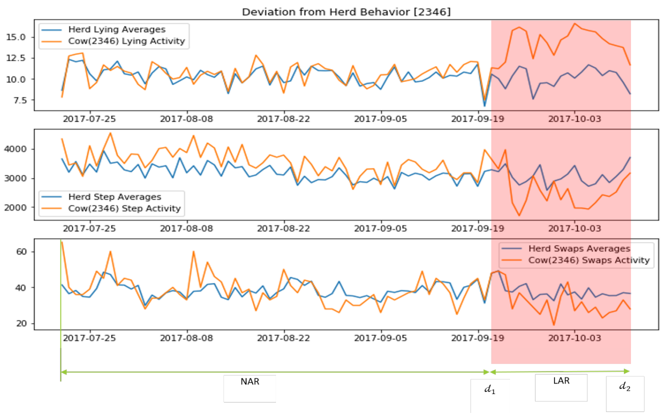

Fig. 4: Relationship between herd mean and cow activity for cow 2346, showing deviation in its behaviour from herd as it transitions into lameness.

same way. Not all the animals in the herd had their activity tracing the herd mean — some had higher, some lower and some equal. This observation led us to our next decision in the analysis, which was to identify the clusters in the herd.

2) **Clustering:** *Does each animal in the herd need to be treated separately i.e., treating each cow as a single experimental unit; or can some clustering technique be used to define clusters of animals that share similar features within the herd?*

The same study [9] referred earlier in forming **Cow Profiles** also shows that cattle in the same pasture are not treated as independent experimental units because of the potential confounding effects of the herd's social interactions. It also provides the insight that activity patterns of groups of cows with in the herd may have level of independence that is sufficient for analysing them as individual units under situations such as large herd size of around 53-240 cows. This means that smaller sizes herd (less than or equal to 40) don't exhibit any patterns of group formations within the herd while the larger herd sizes (53-240) show formations of groups within the herd. It should also be noted that the technology based automated smart solutions for animal welfare are more beneficial for farms with large herd sizes; one can assume that for small ones that the farmer can manually keep track of each animal's welfare without much effort.

From our analysis and literature study, it was clear that a one-size-fits-all approach where it is assumed that all animals behave the same way, and all cows are treated as a single set (i.e. without any grouping) to detect anomaly in behaviour won't be efficient. There exist sub-sets in the herd that share similar features, which once identified can be leveraged to fit the use-case as opposed to a one-size-fits-all solution. In our analysis, we found that even animals of the same age behaved differently and had different levels of activity.

Our clustering model is based on the observation that there were some animals in the herd whose activity levels (step count, lying time and swaps) were always greater than the mean activity value of the herd, while some whose activity levels were always less than the mean herd activity, and then there were others who traced the herd mean. Based on this, we form three clusters as follows:

- **Active:** These are animals in the herd whose activity levels are always higher than the herd mean.
- **Normal:** These are animals in the herd whose activity levels always trace out the herd mean.
- **Dormant:** These are animals in the whose activity levels are always lower than the herd mean.

It's worth mentioning that prior to finalizing activity based clustering in our use-case, we also used age-based clustering [10] to define clusters and then fed those into the classification model for early detection of lameness. It didn't lead to early detection of lameness, and in-line with literature studies we looked for other clustering techniques as well, and found that activity based clustering technique performs better [11] in the use-case of early detection of lameness.

The above conclusion led to further investigation of the clusters, concerning their nature as static clusters, re-clustering, and optimal approaches to clustering. From our analysis, we found that clusters are dynamic in nature, i.e., the animals can migrate from one cluster to another in a time window. There can be a number of reasons behind this, we postulate because of age and weather at least, and perhaps other factors that affect the activity levels of the animals and the herd as a whole.

Thus, it is the responsibility of the clustering model to re-cluster the animals prior to feeding it into the classification model. The optimal time to re-cluster was found out to be about 2 weeks (14 days). This decision was made by continuously observing the movement of animals between different clusters, and finding the time frame of these movements.

3) **Classification — Early Lameness Detection:** *The next important question was to decide on which classification model to use given the objective of early prediction of lameness in dairy cattle?*

Classification algorithms belong to the set of machine learning algorithms that output a discrete value. Often, these output variables are referred to as labels, classes or categories. Classification problems with two classes are called binary classification problems, and those with more are referred to as multi-class. In our use-case scenario, the problem was written as a binary classification problem with Lame being the positive class and Non-lame as the negative class. The data split was as 80-20, i.e. 80% of data was used for model training and rest 20% was used for testing.

We examined a number of classification algorithms [12] ranging from Support Vector Machine (SVM), Random Forest (RF), K-Nearest Neighbors (K-NN) and decision trees. We found that the K-NN based classification algorithm served best for early lameness prediction in our use case scenario, as it was best balanced in terms of accuracy and early lameness detection window. It gave an accuracy of 87% with 3 days early prediction window in advance of any visual sign of lameness to be observed by the farmer.

A short demo video of the overall end-to-end IoT solution thus designed and developed is available at [13].

VI. Benefits to Stakeholders

The detailed impact and benefits to stakeholders are outlined below:

- ***Animals:*** Animals can't communicate the way humans do. With a little bit of technology, we can understand their natural behaviour and trends. We can see the irregularity and change in their behaviour and can then take appropriate measures towards their well being. This not only helps improving the production capacity, but also improves the health and social interactions within the herd.
- ***Farmer:*** Increased size and scale of the farm poses various challenges for a farmer. In this tech-savvy and data-driven era, it's easier for a farmer to manage the well-being of a big herd on a handheld digital device.

VII. Conclusion

We have outlined the key design principles used in the development of our IoT solution aimed at early detection of lameness in dairy cattle. We present the critical decisions made and methodologies used in designing an end-to-end software system in fog-enabled IoT scenarios for our use-case.

The key takeaways are as below:

- A hybrid machine learning model such as one presented — activity based clustering combined with classification model — returns accurate results in detection of anomalies in animal behaviour for early detection of lameness as opposed to one-size-fits-all approach.
- Results clearly suggest that once monitored, the behavioural changes when animals are ill can be mapped to specific illnesses such as lameness in our use-case scenario.
- Many of these behavioural changes that occur before visual onset are extremely subtle, and difficult to detect in practice without technology.
- A careful coordination of computational resources along the technology path from sensor to cloud continuum is vital to the performance of such a system. Edge, fog and cloud resources each bring their unique input towards the functionality and performance of the overall IoT application system developed.

We believe that the insights from this study can contribute to the behavioural analysis of animals, and can help detect subtle changes in livestock behaviour before any clinical symptoms of disease are visible. This will lead to improved insights in animal behavioural analysis, and better practices for farmers. The wearable technology for livestock in conjunction with advanced machine learning methods has the potential for development of robust early warning systems to detect disease development early on.

VIII. Ongoing and Future Work

To further validate the proposed approach for early lameness detection, we are expanding the work undertaken to date through the execution of a use case in the IoF2020 project[1] named MELD[2]. The MELD project will build and expand upon this existing work, and integrate it into the IoF2020 dairy farming technology trials with planned deployments in Portugal, Israel and South Africa, leveraging sensor technologies from two different vendors on a combined total of approximately 1000 cattle. With more data at hand, we then aim to examine other possible clustering techniques and evaluate other classification techniques to further improve the algorithm.

[1]Internet of Food & Farm 2020, https://www.iof2020.eu/

[2]MELD stands for **M**achine **L**earning based **E**arly **L**ameness **D**etection in Beef and Dairy Cattle

Acknowledgments

This work has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) and is co-funded under the European Regional Development Fund under Grant Number 13/RC/2077. Mohit Taneja is also supported by CISCO Research Gift Fund.
The future work (MELD) is funded through the IoF2020 which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 731884.

References

[1] "Eu agricultural market briefs - world food consumption patterns – trends and drivers," https://ec.europa.eu/agriculture/sites/agriculture/files/markets-and-prices/market-briefs/pdf/06_en.pdf, June 2015, (Accessed on 12/14/2017).

[2] "The food and agriculture organization (fao) of the united nations - dairy production and products: gateway to dairy production and products," http://www.fao.org/dairy-production-products/en/#.WdTbDmhSyUl, (Accessed on 12/14/2017).

[3] H. Kharas, "The emerging middle class in developing countries," *OECD Development Centre Working Papers*, 2010. [Online]. Available: http://dx.doi.org/10.1787/5kmmp8lncrns-en

[4] A. Van Nuffel, I. Zwertvaegher, L. Pluym, S. Van Weyenberg, V. M. Thorup, M. Pastell, B. Sonck, and W. Saeys, "Lameness detection in dairy cows: Part 1. how to distinguish between non-lame and lame cows based on differences in locomotion or behavior," *Animals*, vol. 5, no. 3, pp. 838–860, 2015. [Online]. Available: http://www.mdpi.com/2076-2615/5/3/387

[5] "Economic cost of lameness in irish dairy herds," https://www.xlvets.ie/sites/xlvets.ie/files/press-article-files/XLVets%2520Article%2520Forage%2520Guide%25202012.pdf, (Accessed on 05/07/2019).

[6] M. Taneja, N. Jalodia, J. Byabazaire, A. Davy, and C. Olariu, "Smartherd management: A microservices-based fog computing–assisted iot platform towards data-driven smart dairy farming," *Software: Practice and Experience*, vol. 49, no. 7, pp. 1055–1078, 2019. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/spe.2704

[7] "Getting to know MQTT," https://www.ibm.com/developerworks/library/iot-mqtt-why-good-for-iot/index.html, (Last accessed on August 03,2017).

[8] B. Butzin, F. Golatowski, and D. Timmermann, "Microservices approach for the internet of things," in *2016 IEEE 21st International Conference on Emerging Technologies and Factory Automation (ETFA)*, Sept 2016, pp. 1–6.

[9] M. B. Stephenson and D. W. Bailey, "Do movement patterns of gps-tracked cattle on extensive rangelands suggest independence among individuals?" *Agriculture*, vol. 7, no. 7, 2017. [Online]. Available: http://www.mdpi.com/2077-0472/7/7/58

[10] M. Taneja, J. Byabazaire, A. Davy, and C. Olariu, "Fog assisted application support for animal behaviour analysis and health monitoring in dairy farming," in *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*, Feb 2018, pp. 819–824.

[11] J. Byabazaire, C. Olariu, M. Taneja, and A. Davy, "Lameness detection as a service: Application of machine learning to an internet of cattle," in *2019 16th IEEE Annual Consumer Communications Networking Conference (CCNC)*, Jan 2019, pp. 1–6.

[12] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[13] M. Taneja, N. Jalodia, J. Byabazaire, A. Davy, and C. Olariu, "Smartherd – connected_cows_demo.mp4 – google drive," https://drive.google.com/file/d/1QIrKSp8SkAZRRAFQDQHdPXu1TVYaVyv3/view, (Shared on Google Drive).

Mohit Taneja is currently working as a software research engineer, and also a PhD researcher with Emerging Networks Laboratory at Telecommunications Software and Systems Group of Waterford Institute of Technology, Ireland. He is a part of the SFI-CONNECT-IBM project termed SmartHerd, which has now been extended with follow-on IoF2020 project MELD, which is a project to detect early stage lameness in cattle with the help of technology. His research focuses on fog computing support for Internet of Things applications. He received his Bachelor's Degree in Computer Science and Engineering from The LNM Institute of Information Technology, Jaipur, India in 2015.

Nikita Jalodia is a PhD Researcher in the Department of Computing and Mathematics at the Emerging Networks Lab Research Unit in Telecommunications Software and Systems Group, Waterford Institute of Technology, Ireland. She is working as a part of the Science Foundation Ireland funded CONNECT Research Centre for Future Networks and Communications, and her research is based in Deep Learning and Neural Networks, NFV, Fog Computing and IoT. She received her Bachelor's Degree in Computer Science and Engineering from The LNM Institute of Information Technology, Jaipur, India in 2017, with an additional diploma specialization in Big Data and Analytics with IBM.

Paul Malone graduated from Waterford Institute of Technology with a 1st class honours degree in 1998 and completed a research MSc in 2001. Paul has worked on researching technologies and techniques related to IT security in the areas of distributed trust and reputation management and privacy and data protection controls. Paul is currently coordinating an Agri-Tech use case sub-project of the H2020 IoF2020.eu project applying machine learning technologies in detecting early stage lameness in cattle.

John Byabazaire is currently a PhD student in school of computer science, University College Dublin working on IoT systems for data collection in precision agriculture. Before that, John pursed a MSc in Computer Science from Waterford Institute of Technology and due to graduate. He received his BSc in Computer Science from Gulu University in 2013. Alongside his current PhD study, John continues to research e-learning for low bandwidth environments, Software Defined Networking, Network Function Virtualization, and Remote Sensing, IoT (Internet of Things) and Fog Analytics.

Alan Davy received the B.Sc. (with Hons.) degree in applied computing and the Ph.D. degree from the Waterford Institute of Technology, Waterford, Ireland, in 2002 and 2008, respectively. He is currently Head of Department of Computing and Mathematics in School of Science and Computing Waterford Institute of Technology (WIT), Waterford, Ireland. Previously, he was Research Unit Manager of the Emerging Networks Laboratory with the Telecommunications Software Systems Group of WIT. He is coordinator of a number of national and EU projects such as TERAPOD. His current research interests include Virtualised Telecom Networks, Fog and Cloud Computing, Molecular Communications and TeraHertz Communication.

Cristian Olariu received his B.Eng. degree in 2008 from the Faculty of Electronics and Telecommunications which is part of Politehnica University of Timisoara, Romania, and his Ph.D. degree in 2013 from the Waterford Institute of Technology, Ireland, in VoIP over wireless and cellular networks. He is currently a research engineer with the Innovation Exchange, IBM Ireland, and previously he was a post-doctoral research fellow with University College Dublin, Ireland. His interests are in the field of Automotive Services, IoT, Vehicular Communications, Voice over IP, QoS provisioning for Voice over IP, Wireless Networks, Software Defined Networking and Network Functions Virtualisation.