Avena: An Event-Driven Software Framework for Informed Decisions and Actions in Cropping Systems

Fabio A. Castiblanco[1], Md. Samiul Basir[2], Andrew D. Balmos[1], James V. Krogmeier[1], Dennis R. Buckmaster[2]

[1] Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, 47907, U.S.A

[2] School of Agricultural and Biological Engineering, Purdue University, West Lafayette, IN, 47907, U.S.A

A paper from the Proceedings of the
16th International Conference on Precision Agriculture
21-24 July 2024
Manhattan, Kansas, United States

Abstract

Interoperability is one of the enabling factors of real-time communications and data exchange between asynchronous data actors. Interoperability can be attained by introducing the notion of events to systems that extract data from consumed ground-truth event streams that utilize application-specific structures. Events are specific occurrences happening at a particular time and place. Event-data are observations of phenomena, or actions, as seen by different systems in Internet of Things (IoT) deployments, independent from data structures and platforms. Avena, an open-source software framework, facilitates communication among systems ranging from IoT devices, agricultural machines, and cloud-based data processing services and fosters effortless interoperability. Avena addresses the challenges of intermittent connectivity and incompatibility among precision agricultural systems, facilitating resilient IoT communications in complex agricultural networking environments. Constructed on modern distributed system architectures and leveraging overlay networks based on WireGuard tunnels and the NATS messaging system, Avena simplifies connections among software and other systems, making them inherently tolerant to communication outages and seamlessly adaptable to the dynamic nature of agriculture. Device discovery mechanisms and peer-to-peer communications between devices in precision agriculture ecosystems are integral components of Avena with event-level data abstraction and data interoperability. In this work, we integrated agricultural operational data which is mostly textual and manually recorded by farm employees. We have integrated Meta Ag, an open-source Android application featuring a geofence-responsive information bot that makes the in-field record-keeping process automatic and less error prone. Meta Ag, integrated into the Avena platform and powered by NATS, can give useful insights into machine and sensor data providing complete in-field operational records. With machine and sensor data now interoperable with contextual data, data-driven decision-making, including integration with biophysical models, will be easier, secure, and reliable.

Keywords.
Distributed system, IoT, Data Pipeline, Avena, Precision Agriculture

Introduction

With the advancement of agricultural systems from mechanization to automation, communications has become a priority. Communication in agriculture has evolved from cellular (Tang et al., 2021) to wireless sensor networks (Mohd Kassim et al., 2014) and Internet of Things (IoT) (Liang & Shah, 2023) to facilitate data driven decision support and real time data transmission.

IoT communications comprise the exchange of data among various agricultural things (sensors, computers, controllers, etc.), leveraging data from a wide array of sources and structures. Farm machinery, equipped with advanced data loggers, generates comprehensive machine data, including fuel usage, machinery position, temporal information, oil and tire pressure, hydraulic data and many more (B. McLaughlin et al., 1993). In-field sensors provide site-specific sensory data such as soil moisture levels, temperature, Leaf Area Index (LAI), wind speed, precipitation, etc. (Li et al., 2010). Furthermore, UAVs (Unmanned Aerial Vehicles) and satellites capture high-resolution imagery data, enriched with spatial and temporal attributes (Brown, 2015; Delavarpour et al., 2021; Tsouros et al., 2019; Zhang et al., 2021). These images are processed to derive critical metrics such as Normalized Difference Vegetation Index (NDVI), LAI, and Normalized Difference Water Index (NDWI) using sophisticated analytical models (Xue & Su, 2017). This diverse and complex data is seamlessly transmitted among various entities within the IoT system, enhancing the overall efficiency and decision-making processes in agriculture.

In addition to the aforementioned data sources, operational records play a crucial role in agriculture (Baumüller & Kah, 2019; Daum et al., 2021; Fabregas et al., 2019). These records are essential for providing context to the data generated by sensors, UAVs, and machinery. Detailed activity records, represented as descriptive metadata, offer valuable insights that help to interpret and understand the raw data collected from various sources (Zinke-Wehlmann et al., 2021). This contextual metadata is generally collected through manual inputs, making it predominantly textual in nature. The applications involved in gathering this metadata are distinct from other data sources within the IoT system. Despite these differences, for the contextual metadata to be effective and useful, it needs to be integrated into the same distributed system that handles sensor, machinery, and remote sensing data. By incorporating contextual metadata into the data driven system, it becomes possible to create a comprehensive and cohesive dataset. This integration enhances the ability to analyze and utilize the data effectively, leading to more informed decision-making and improved outcomes in agricultural operations.

To account for multiple things operating autonomously within a single IoT environment we must utilize distributed computing techniques (El-Sayed et al., 2018). Distributed computing systems consist of multiple interconnected computers that work to perform complex tasks or applications by dividing the workload among themselves. These systems leverage the combined processing power, memory, and storage capabilities of various machines, typically located in different geographical locations, to enhance performance, increase scalability, and ensure fault tolerance. Common applications include data processing, scientific simulations, and large-scale web services, which benefit from distributed architectures by achieving higher efficiency and reliability compared to using a centralized computing system (Hwang et al., 2011).

The diverse nature of data sources and structures is a primary challenge to achieving interoperability. In a distributed system, all computing nodes need to communicate with each other, even those not initially envisioned. Data shared from one source may be in an unfamiliar or proprietary format, creating barriers for other nodes to utilize it. To address data interoperability, we propose a software framework based on Event-driven Architectures (EDA). In this approach, data is transferred as a single message unit representing an observation of an event, independent of a predefined structure.

Events are the fundamental units of EDAs and depend on time and location. They can be portrayed as a change in the state of a service, such as the beginning or finalization of an activity (Laliwala & Chaudhary, 2008). EDAs focus on the comprehension of an event at its primary abstraction (McGovern et al., 2006). Being asynchronous, this communication architecture becomes an integral part of a distributed computing system which handles the co-occurrence of events in multiple entities within a single system without requiring awareness of the state of other distributed entities within the system (Kshemkalyani & Singhal, 2011). Events as data transmission units in real-time data streams of IoT environments bring the solution of interoperability between different data sources and stakeholders. Event messages not only solve data format issues, but also address data granularity in spatial and temporal aspects. Event data from machine sensor systems of a tractor have different spatial granularity from UAV generated data. Similarly, field operational metadata is recorded for the entire duration of a field operation. Addressing this granularity within IoT environments becomes challenging.

This work proposes an approach to integrate diverse data streams to synergize real-time in-field data. Specifically, we combine user input data on agricultural operations with IoT sensor streams. The goal is to contextualize event data streams using manual records to enhance decision support, demonstrating the integration of multiple data sources and granularities. The specific objectives of this work are:

- Integrate automated data streams from IoT systems with user applications through an event-driven software framework.
- Connect disparate data sources in precision agriculture systems using network abstraction and distributed computing services.

System Overview

The Avena Software Framework

Avena is a software framework designed for data interoperability among disparate data sources and computational instances. Built as a collection of open-source technologies, Avena integrates data flows at various levels. Ranging from sensing devices in precision agriculture scenarios, to remote computing instances running algorithms for enhanced decision-making in agricultural functions. At the foundational level, Avena uses the GNU/Linux operating system across both edge-computing nodes and fixed/cloud servers. It manages services running on these platforms that accomplish tasks that range from abstracting hardware functions, accomplishing local computing needs, provides interfaces to access data not just from local sensors and actuators, but also other remote Avena based devices.

To accomplish this, while also retaining security and multi-tenancy, Avena utilizes Open Container Initiative (OCI) software containers to provide process and data isolation. Containers also offer easier software packaging and distribution as they are packaged as a single unit that includes all necessary software dependencies. Regardless of the computing platform in use, the application should still run.

At the network layer, WireGuard is used to create a Virtual Private Network (VPN) between all cooperative Avena devices. WireGuard is a simple and secure data transmission protocol that abstracts all the intermediate network links between data sources and consumers.

Lastly, at the application layer, the messaging system NATS facilitates communications between logical assets. NATS provides interest-based message routing, in which software services running at different stages in systems share data flows between them via Avena event messages (Balmos et al., 2022).

In the Avena ecosystem, each component is designed to be independent and isolated from a particular version or instance of another thing. Instead they are dependent the logical representation of what data those things may need and/or provide. This interface is created by a standard set of message definitions that try to be as minimal and as atomic as possible, and therefore, as interoperable as possible. To achieve this, three primary constraints apply:

1. Devices need to communicate without knowing about each other: The network should

deliver messages to their logical destination as opposed to address a specific piece of hardware.

2. Data needs to be as atomic as possible: Components should only have to understand the data they need, not the device or thing that created it. In other words, event messages should be as small and as specific as possible. For example, instead of a message which includes all of machines status information, there should be many messages in which each status is communicated separately. Then a downstream consumer only interested in machine location must understand location events, as opposed to a complex full machine status message which is specific to a certain type of machine.
3. Information needs to be addressed based on interest. Data producers should not choose where its events are sent to, instead a receiving device should indicate to the network which type of messages it is interested in, and the network works to route those messages to it.

The components of the Avena ecosystem that accomplish this are described below.

Message Queues

Message queues are data exchange mechanisms used in computing to facilitate asynchronous communication between different processes or systems. They operate by temporarily storing messages in a queue until they can be processed by the receiving application. This method helps decouple components by allowing them to operate independently, managing fluctuations in workload and ensuring data is not lost when the receiving system is busy or down. Message queues are essential for enhancing scalability, reliability, and fault tolerance in distributed systems. Popular implementations include NATS, RabbitMQ, Apache Kafka, and AWS SQS.

Overlay Networking

Overlay networks are an approach of network abstraction to create a simplified network layer on top of the existing physical network infrastructure. This technique allows multiple virtual networks to coexist and operate independently over the same hardware, enabling flexible data routing, improved network management, and enhanced security features without altering the underlying physical network setup. Overlay networks are commonly used in cloud computing, data center architectures, and to support virtual private networks (VPNs).

Virtual Private Networks (VPN) is a technology that creates a secure, encrypted connection over less secure network links, such as the Internet. VPNs are used to establish secure connections, encrypting internet traffic and online identity. This technology serves as a network abstraction mechanism by creating a secure and private network connection over a public infrastructure like the Internet. This abstraction allows users and devices to exchange data as though they were directly connected to a private network, irrespective of their actual, physical network connections. Since VPNs encapsulate and encrypt data packets, they ensure that data transmitted between devices remains secure and private, even when sent over insecure or public networks.

In addition to abstractions at the data-link layer via VPNs, Avena implements Interest-based routing which is a network communication method where data packets are transmitted based on the interests or data requirements expressed by recipients, rather than using conventional IP address-based routing. This model is used in publish-subscribe systems and content-centric networks, where nodes specify the type of content they wish to receive, rather than specifying particular senders. This approach optimizes data distribution by ensuring that nodes only receive information that is relevant to their declared interests, reducing unnecessary data transmission and enhancing network efficiency in highly intermittent connectivity scenarios (Ayu et al., 2019). Interest-based routing is particularly useful in scenarios with high data volume and dynamic content distribution, such as streaming media, distributed sensor networks, and large-scale IoT deployments.

Edge computing, IoT and Sensors

Edge computing is a distributed computing paradigm that brings computation and data storage closer to the location where data is generated, to improve response times and save bandwidth. This approach involves processing data near the edge of the network, close to where it is generated, rather than at a centralized data center. Edge computing is particularly beneficial in scenarios where low latency is critical or where bandwidth is limited, such as in IoT devices, mobile computing, and autonomous vehicles.

Internet of Things (IoT) are networks of interconnected physical devices equipped with sensors, software, and other technologies to connect and exchange data with other devices and systems over the internet. IoT extends internet connectivity beyond traditional devices like computers and smartphones to a diverse range of everyday objects and industrial equipment, enhancing automation, real-time monitoring, and data collection.

Sensors are devices or instruments that detect and respond to various types of physical inputs from the environment, such as light, heat, motion, moisture, pressure, or any other entity. They convert these inputs into readable signals, typically electrical, which can be measured, transmitted, and analyzed, allowing for more informed decision-making and automation across various applications.

Abstraction Services

According to Ziegler & Dittrich (2007), the problem of integration in information systems relies on combining systems in a way that they form a unified whole and provide users with the experience of interacting with a single information system. In Avena, abstraction is the approach to build simpler interactions between software components in digital agriculture ecosystems. By focusing only on necessary interactions and inputs, software designs rely on shared data schemas defined by the needs of the members of the ecosystem, such as data processing services, data visualization platforms, IoT systems, decision-making support tools, among others. Abstraction allows interoperability by establishing a common framework where different systems communicate and work together regardless of their underlying architectural complexities. In the same way, hiding technology-specific complexities allows simpler interactions between the components of a system, making it easier to understand the overall functioning and reducing the overhead in data exchanges. To achieve integration, Avena implements two types of data abstractions depending on the level of abstraction, and the use case of a particular data flow.

Hardware Abstraction Services

A Hardware Abstraction Service (HAS) is software that bridges hardware data to the Avena ecosystem by translating a data stream from a device, often in a vendor-specific format, into a standardized message specification. These software services are typically written by developers with extensive knowledge on how a device operates and are also familiar with the Avena message specification. Consequently, a HAS publishes hardware data streams to the messaging system, enabling other applications and services to access device data without needing to know specific details such as hardware configurations or underlying protocols (Balmos et al., 2022). In many cases, there are multiple providers of hardware that accomplish the same or similar goals, in fact it is exactly this case that makes HASs so important. Each provider would have a HAS that is unique to their product while exposing the same standard interface as the others. In that way, downstream services which want to leverage the hardware do not have to know about product specific features, functionalities, or workaround. They just rely on a properly built HAS service to make the necessary translations. The typical data flow from HAS is shown in Figure 1.

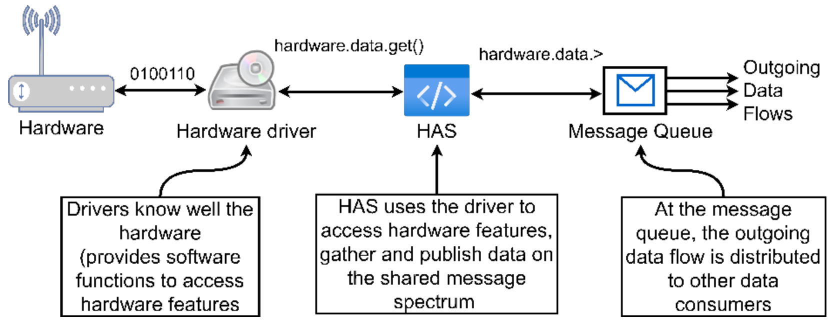

Figure 1. Hardware Abstraction Service.

Content Abstraction Services

A Content Abstraction Service (CAS) plays a role in the Avena ecosystem similar to that of a HAS, in the sense they are able to serve as data sources. However, unlike a HAS, which adapts hardware for integration into the Avena ecosystem, a CAS merges data flows from various services to provide contextually enhanced output. For example, hardware data flows can be enhanced by incorporating information from other data streams through computations that correlate fields between them. As a result, the combined data streams can address more complex questions, such as 'How much work was done in a place?' by synthesizing information from diverse sources (Balmos et al., 2022). Data received and publishing procedure of a CAS is shown in Figure 2.

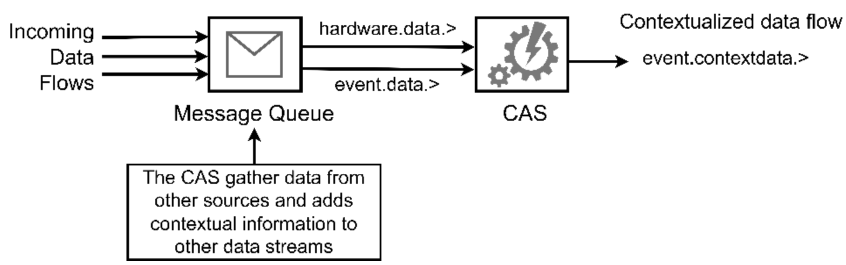

Figure 2. Content Abstraction Service.

Computational Models

Computational models are mathematical models that represent complex systems or phenomena through algorithms, using the capabilities of a computer to simulate and study their behavior. They are typically used to predict future events based on specific parameters or to analyze scenarios that are too complicated or costly to recreate. Common applications include weather forecasting, economic simulations, and biological processes. Models help scientists and researchers explore and understand complex interactions and dynamics efficiently and accurately.

Storage

In Avena, data storage is categorized into short-term and long-term storage based on the timeframe of the application. Short-term storage is utilized for data that requires active access or modification, such as in transaction processing systems, temporary caches, or during session-based computing where quick accessibility is crucial. Data in this layer maintains the same message schema as when transferred, allowing for an accurate recreation of event data.

Long-term storage, on the other hand, is dedicated to preserving data for extended periods, often for archival, compliance, historical, or backup purposes. This type of storage is designed to be both durable and reliable, ensuring that data remains accessible and intact over years or even decades. Moreover, data streams from short-term storage can be transformed into long-term storage batches, depending on their application. For example, a single short-term data stream may generate one or more long-term data batches, with each batch tailored to a specific use case and potentially employing different schemas adapted to the application's requirements.

Short-term storage is often implemented through storage caches or persistence layers in messaging systems, while long-term storage can be implemented through database management systems.

Meta Ag

Meta Ag (Basir et al., 2023) is an open-source farm contextual metadata recording application developed for Android. It automatically captures users' information, location and time of a farm activity using GPS location. Further precise details about specific activities are also captured with the help of a rules-based chatbot mitigating errors inherent in traditional data collection methods through data validation options. Meta Ag minimizes missing and inaccurate data of field activities. Moreover, it advocates for data interoperability and open-source technology, making it suitable for research and development. Meta Ag can function independently or seamlessly integrate with other Farm Management Information Systems (FMIS) for comprehensive agricultural activity records.

Working Principle of Meta Ag

Meta Ag works with six major components, namely, Authentication, Geofence creation, Geofence detection, Data Collection Wizard, Setting Infobot Options, and Data Access Authentication allows only the registered users to use the system. It provides data security as well as collecting user information. The Geofence creation module helps mark the fields in a farm or farmer working fields, which are to be detected by the app. The geofence detection module collects users' location in the background and matches the location with the created field boundaries if the user's location falls inside any boundary. This module collects the entry and exit time of user to and from a field, stores the information and informs the user about the field activity. Data Collection wizard asks the user for further details about the accessed field. It uses a rules-based chatbot, Infobot, to collect the details on what activity was conducted, how, and to what extent. The Setting Infobot Option module helps set the data validation options given by Infobot. The options in Infobot are clickable and user do not require to type most of the entries. This Infobot helps the app to record data without any missing or discrepancy. The collected data is stored in Google Firebase database. The Data access module allows authentic users to download the data in csv format.

Metadata Generated by Meta Ag

The metadata produced by Meta Ag comprises JSON string data that encompasses various details pertaining to in-field activities. The data recorded by Meta Ag encompasses the following aspects:

- Date, time, and duration of the activity.
- Information regarding the individuals involved.
- Name and description of the activity.
- Details regarding planting/sowing operations.
- Involvement of machinery and tools, along with pertinent information.
- Scope and extent of the undertaken work.
- Utilized materials specific to the field activity.
- Additional observations and relevant notes.

A visual representation illustrating the collected activity details by Meta Ag is depicted in Figure 3.

```
"Operation": {
  "9CS": {
    "1654341660000": {
      "Sowing": {
        "activity_type": "Line Sowing",
        "asset_description": "6400 Cab",
        "asset_note": "leftover seed ",
        "asset_used": "Cab",
        "crop": "Corn-P0720AM",
        "croprate": "36000",
        "duration": "3Hr. 25mins. 0sec.",
        "entry": "2022-06-04 07.21.00",
        "exit": "2022-06-04 10.46.00",
        "implement": "JD 7000",
        "material": "Nothing",
        "material_amount": "",
        "material_type": "",
        "operator": "Samiul",
        "work_amount": "leftover seed was planted"
      }
    }
  },
```


Figure 3. Meta Ag generated activity detail in JSON format.

Integration of MetaAg data to an Avena ecosystem

Meta Ag functions as an automated metadata generator within the Avena platform, replicating the role typically performed by an Avena user by publishing metadata within the Avena framework. To publish metadata in Avena, a content abstraction service is required for metadata extraction and abstraction. The metadata abstraction service within the app retrieves metadata from the app and decomposes it into either discrete message for each record or aggregates it into comprehensive metadata for a single operational event. These atomic messages are then published to a message queue, making them accessible to data consumers within the Avena ecosystem through a subscription model.

Data consumers contextualize the data they subscribe to from various sources or generate for specific events and subsequently publish the contextualized data back to Avena. Subscribers to event messages within the Avena system can thus achieve a comprehensive and contextualized understanding of event data. This process ensures that all stakeholders have access to enriched, context-aware data that enhances decision-making and operational efficiency within the Avena ecosystem.

Case Study: Contextualize CANBUS Data for Agricultural Operations

In this scenario, a corn harvester equipped with an ISOBlue device is used for real-time CANBUS data logging while harvesting a field. The harvester publishes its real-time location data within the Avena environment. The automatic flow of CANBUS data can provide temporal and spatial information about the vehicle but lacks contextual details such as the specific activity being performed, the quantity harvested, the operator's identity, the field name, and any observations made by the operator.

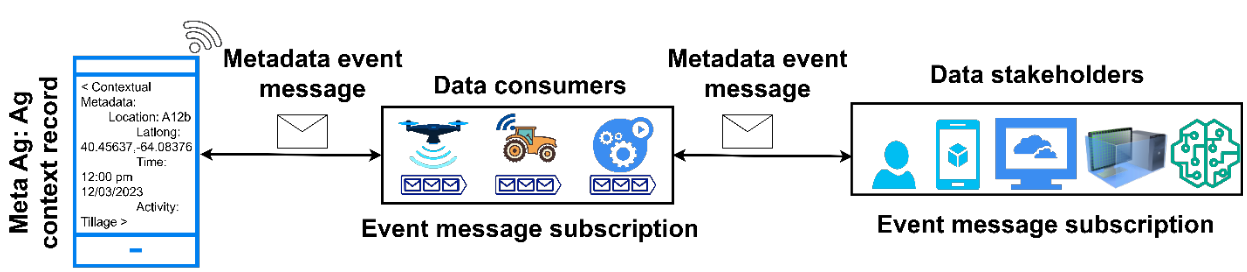

Figure 4. Schema of Meta Ag communication in Avena environment.

To address this, the operator employs the Meta Ag app to record detailed contextual metadata about the harvesting event, which is then published in the Avena environment. The ISOBlue device capturing the CANBUS data also subscribes to this metadata. As Meta Ag updates the system post-operation, Avena utilizes two Content Abstraction Services (CASs): one for processing the CANBUS data and another for processing the contextual metadata. These services integrate data from both sources to construct a comprehensive, contextualized CANBUS dataset for the harvesting operation. Figure 5 shows the contextualization strategy of CANBUS data in Avena framework for a holistic operational event information.

The resultant dataset, which is stored in a cloud environment, includes detailed information such as vehicle health, fuel usage, operation date and location, operator identity, and specific harvesting observations. This enriched dataset can be leveraged for various operational optimizations, including maintenance scheduling, fuel refills, operator assignments, harvesting decisions, and site selection.

By integrating real-time CANBUS data with contextual metadata through the Avena framework and Meta Ag, a holistic view of vehicle operations is achieved, enhancing both the efficiency and effectiveness of agricultural practices.

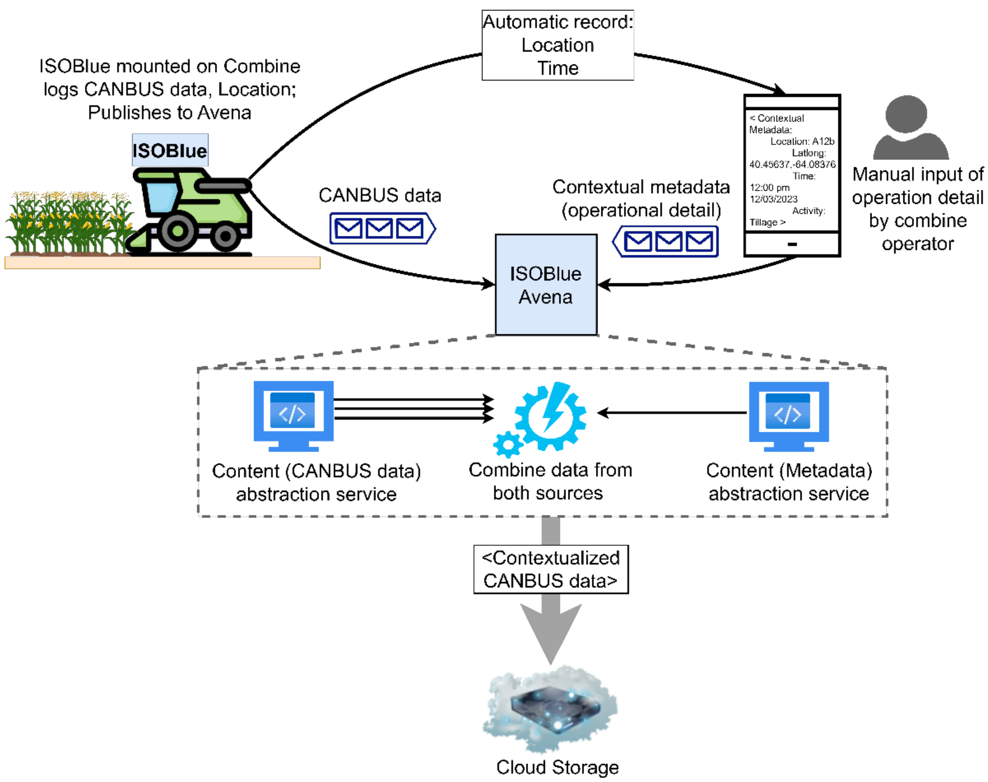

Figure 5. Contextualizing Vehicle Operations with Meta Ag in the Avena Framework.

Summary

Avena, an open-source framework for IoT communication and computing in agriculture, enables seamless and interoperable real-time data streaming among various data stakeholders. Operating in a publish-subscribe mode, Avena exchanges data in an atomic message format that follows predefined schema or specification, thereby enabling interoperability between components. The framework handles diverse data types, including sensor, machinery, UAV data, and user inputs such as contextual metadata.

In this article, we discussed a use case where Avena manages two different types of data from various sources, specifically illustrating the integration of the Meta Ag app with Avena. This example highlights Avena's capability to process data with varying spatial and temporal attributes, demonstrating its versatility and robustness. Furthermore, the use case exemplifies the addition of application-generated data into the Avena framework, showcasing its ability to enhance data contextualization and enrich the overall dataset.

Avena facilitates efficient data management and enables data interoperability, which is crucial for informed decision-making in agriculture. By supporting real-time data streaming and contextual metadata integration, this framework significantly improves operational efficiency, resource management, and overall productivity. This showcases Avena as a tool for advanced precision agriculture and promoting sustainable farming practices. Future research could explore expanding Avena's capabilities to integrate additional data sources and enhance predictive analytics, further contributing to the evolution of smart agriculture.

Acknowledgments

This work is supported by the IoT4Ag Engineering Research Center funded by the National Science Foundation (NSF) under NSF Cooperative Agreement Number EEC-1941529. It was also funded in part by the NSF under CNS-2235134. Any opinions, findings and conclusions, or recommendations expressed in this material are those of the author(s), and do not necessarily reflect those of the NSF.

References

Ayu, V., Soelistijanto, B., & Sijabat, J. (2019). Interest-based Epidemic Routing in Opportunistic Mobile Networks. *2019 7th International Conference on Information and Communication Technology (ICoICT)*, 1–6. https://doi.org/10.1109/ICoICT.2019.8835355

B. McLaughlin, N., C. Heslop, L., J. Buckley, D., R. St. Amour, G., A. Compton, B., M. Jones, A., & Van Bodegom, P. (1993). A General Purpose Tractor Instrumentation and Data Logging System. *Transactions of the ASAE*, *36*(2), 265–273. https://doi.org/10.13031/2013.28334

Balmos, A. D., Castiblanco, F. A., Neustedter, A. J., Krogmeier, J. V., & Buckmaster, D. R. (2022). ISOBlue Avena: A Framework for Agricultural Edge Computing and Data Sovereignty. *IEEE Micro*, *42*(1), 78–86. https://doi.org/10.1109/MM.2021.3134830

Basir, M. S., Zhang, Y., Buckmaster, D. R., Raturi, A., & Krogmeier, J. V. (2023). Meta Ag: An Automatic Contextual Agricultural Metadata Collection App. *2023 ASABE Annual International Meeting*, 1.

Baumüller, H., & Kah, M. (2019). *Going digital: Harnessing the power of emerging technologies for the transformation of Southern African agriculture*. https://doi.org/10.4324/9780429401701-23

Brown, M. E. (2015). Satellite Remote Sensing in Agriculture and Food Security Assessment. *Procedia Environmental Sciences*, *29*, 307. https://doi.org/10.1016/j.proenv.2015.07.278

Daum, T., Villalba, R., Anidi, O., Mayienga, S. M., Gupta, S., & Birner, R. (2021). Uber for tractors? Opportunities and challenges of digital tools for tractor hire in India and Nigeria. *World Development*, *144*, 105480. https://doi.org/10.1016/j.worlddev.2021.105480

Delavarpour, N., Koparan, C., Nowatzki, J., Bajwa, S., & Sun, X. (2021). A Technical Study on UAV Characteristics for Precision Agriculture Applications and Associated Practical Challenges. *Remote Sensing*, *13*(6), Article 6. https://doi.org/10.3390/rs13061204

El-Sayed, H., Sankar, S., Prasad, M., Puthal, D., Gupta, A., Mohanty, M., & Lin, C.-T. (2018). Edge of Things: The Big Picture on the Integration of Edge, IoT and the Cloud in a Distributed Computing Environment. *IEEE Access*, *6*, 1706–1717. https://doi.org/10.1109/ACCESS.2017.2780087

Fabregas, R., Kremer, M., & Schilbach, F. (2019). Realizing the potential of digital development: The case of agricultural advice. *Science*, *366*(6471), eaay3038. https://doi.org/10.1126/science.aay3038

Hwang, K., Jack Dongarra, J., & Fox, G. C. (2011). *Distributed and Cloud Computing From Parallel Processing to the Internet of Things* (1st ed.). Elsevier. https://shop.elsevier.com/books/distributed-and-cloud-computing/hwang/978-0-12-385880-1

Kshemkalyani, A. D., & Singhal, M. (2011). *Distributed Computing: Principles, Algorithms, and Systems*. Cambridge University Press.

Laliwala, Z., & Chaudhary, S. (2008). Event-driven Service-Oriented Architecture. *2008 International Conference on Service Systems and Service Management*, 1–6. https://doi.org/10.1109/ICSSSM.2008.4598452

Li, S., Simonian, A., & Chin, B. A. (2010). Sensors for Agriculture and the Food Industry. *The Electrochemical Society Interface*, *19*(4), 41. https://doi.org/10.1149/2.F05104if

Liang, C., & Shah, T. (2023). IoT in Agriculture: The Future of Precision Monitoring and Data-Driven Farming. *Eigenpub Review of Science and Technology*, *7*(1), Article 1. https://studies.eigenpub.com/index.php/erst/article/view/11

McGovern, J., Sims, O., Jain, A., & Little, M. (Eds.). (2006). Event-Driven Architecture. In *Enterprise Service Oriented Architectures: Concepts, Challenges, Recommendations* (pp. 317–355). Springer Netherlands. https://doi.org/10.1007/1-4020-3705-8_8

Mohd Kassim, M. R., Mat, I., & Harun, A. N. (2014). Wireless Sensor Network in precision agriculture application. *2014 International Conference on Computer, Information and Telecommunication Systems (CITS)*, 1–5. https://doi.org/10.1109/CITS.2014.6878963

Tang, Y., Dananjayan, S., Hou, C., Guo, Q., Luo, S., & He, Y. (2021). A survey on the 5G network and its impact on agriculture: Challenges and opportunities. *Computers and Electronics in Agriculture*, *180*, 105895. https://doi.org/10.1016/j.compag.2020.105895

Tsouros, D. C., Bibi, S., & Sarigiannidis, P. G. (2019). A Review on UAV-Based Applications for Precision Agriculture. *Information*, *10*(11), Article 11. https://doi.org/10.3390/info10110349

Xue, J., & Su, B. (2017). Significant Remote Sensing Vegetation Indices: A Review of Developments and Applications. *Journal of Sensors*, *2017*, e1353691. https://doi.org/10.1155/2017/1353691

Zhang, L., Zhang, Z., Luo, Y., Cao, J., Xie, R., & Li, S. (2021). Integrating satellite-derived climatic and vegetation indices to predict smallholder maize yield using deep learning. *Agricultural and Forest Meteorology*, *311*, 108666. https://doi.org/10.1016/j.agrformet.2021.108666

Ziegler, P., & Dittrich, K. R. (2007). Data Integration—Problems, Approaches, and Perspectives. In J. Krogstie, A. L. Opdahl, & S. Brinkkemper (Eds.), *Conceptual Modelling in Information Systems Engineering* (pp. 39–58). Springer. https://doi.org/10.1007/978-3-540-72677-7_3

Zinke-Wehlmann, C., Kirschenbaum, A., Palma, R., Brahma, S., Charvát, K., Charvát Jr, K., & Reznik, T. (2021). Linked Data and Metadata. *Big Data in Bioeconomy: Results from the European DataBio Project*, 79–90.