Chapter 4
Towards a Digital Twin for the Condition Based Maintenance of a Hawk T1A Aircraft

Matt Tipuric, Matthew S. Bonney, Robin Mills, and David J. Wagg

Abstract Condition based maintenance for aircraft has major potential benefits for both reduction of cost and extension of operational hours. As this strategy relies on using live data to update a virtual representation of a system in order to make operational decisions, the relevance of digital twins is clear. This paper discusses the early-stage development of such a digital twin using vibration data collected from an ex-service Hawk T1A aircraft. The primary output of this research – the dataset itself as a benchmark for structural health monitoring – is detailed elsewhere. However, for digital twins to be successfully developed and employed, the holistic development of the entire cyber-physical system must be considered. The practical aspects of methodologies and frameworks for all parts of the system must be developed and refined. This paper discusses a number of these practical implementations within the context of the Hawk data, including sensor choice and placement, data pipelines, virtual models of legacy systems, and platform design. Based on this discussion, recommendations are made, both for immediate improvements of the data collection process and for the longer-term development of a practical operational digital twin.

Keywords digital twin · aircraft · SHM

Introduction

Digital twins are a technology with the potential to be transformational across multiple domains, both within engineering disciplines and more broadly. Such domains include manufacturing, healthcare, asset management, smart cities, and earth sciences. While the multiple definitions of a digital twin have been proposed over the years since its conception, a broadly applicable definition is *a virtual representation of a physical system with the capacity for bidirectional information exchange in relevant time between the digital and physical components.* Digital twins find use both during the design of new assets and across the lifecycle of existing assets - during the *product design phase* and the *asset management phase,* respectively [1].

One common research field that can utilise digital twins in both the product design and asset management phases is structural dynamics. The verification and validation of models using prototypes is vital for the design of critical components and assemblies, while condition monitoring allows for effective maintenance scheduling, life-time extension, and crisis management. None of these areas inherently require the use of digital twins. However, if well implemented, digital twins can act as enablers for all these areas by reducing siloing, simplifying workflows and allowing for live monitoring of assets.

These uses of digital twins make them particularly attractive in the aerospace sector. Aircraft have complex design requirements, with aerodynamic, lightweight geometries and components being exposed to hostile conditions. In addition to this hostile environment, these systems experience long service lives, often more than 30 years, and as a result are expensive to purchase and repair. They have the potential to fail catastrophically and so are required to undergo frequent and rigorous inspection and maintenance. Digital twins have the potential to improve each of these areas, improving the validation of

Matt Tipuric · Robin Mills · David J. Wagg
School of Mechanical, Aerospace and Civil Engineering University of Sheffield, Sheffield, S10 2TN
e-mail: m.tipuric@sheffield.ac.uk; robin.mills@sheffield.ac.uk; david.wagg@sheffield.ac.uk

Matthew S. Bonney
Peirianneg Awyrofod, Aerospace Engineering Prifysgol Abertawe, Swansea University, Abertawe, SA1 8EN
e-mail: m.s.bonney@swansea.ac.uk

Chad Walber et al. (eds.), *Sensors & Instrumentation and Aircraft/Aerospace Testing Techniques, Vol. 8*, Conference Proceedings of the Society for Experimental Mechanics Series, https://doi.org/10.13052/97887-438-0153-5_4

models at the design phases and increasing the efficiency of in-service condition monitoring. However, much work is required before the practical implementation of digital twins for structural health monitoring at scale.

This paper focuses on an ex-service Hawk T1A aircraft, which underwent dynamic testing in the Laboratory for Verification and Validation (LVV) in the University of Sheffield. The primary focus of this work was the generation of high quality, publicly available data sets to support the future development of condition monitoring techniques. The data sets themselves - both of the starboard wing [2] and of the whole structure [3] - are described in detail in other papers which will not be duplicated here, beyond the small amount which is required for comprehension. A secondary goal of this work was to capture the difficulties in collecting and working with datasets at this scale, in support of the practical development of digital twins. These ancillary parts of the process are the focus of this current paper.

The Structure and Models

The BAE Systems Hawk T1A aircraft (‘the Hawk’) is a two-seater aircraft which was previously used by the British Royal Air Force and was previously used for advanced pilot training. After its retirement from service, it was donated to the University of Sheffield by the Defence Science and Technology Laboratory and is now housed at the LVV. It is 12.4m (40.7 feet) in length, with a wingspan of 9.1m (29.9 feet).

For safety and security reasons, a number of components were removed before the Hawk’s arrival on site. These include the cockpit canopy, seats and instrumentation, the engines, and all electronics. Additionally a number of panels are missing, including some of the port wing, the flaps and rudder, and a substantial amount of the upper, rear fuselage, as shown in Fig. 1. The Hawk rests on its wheels on the strong floor of the LVV.

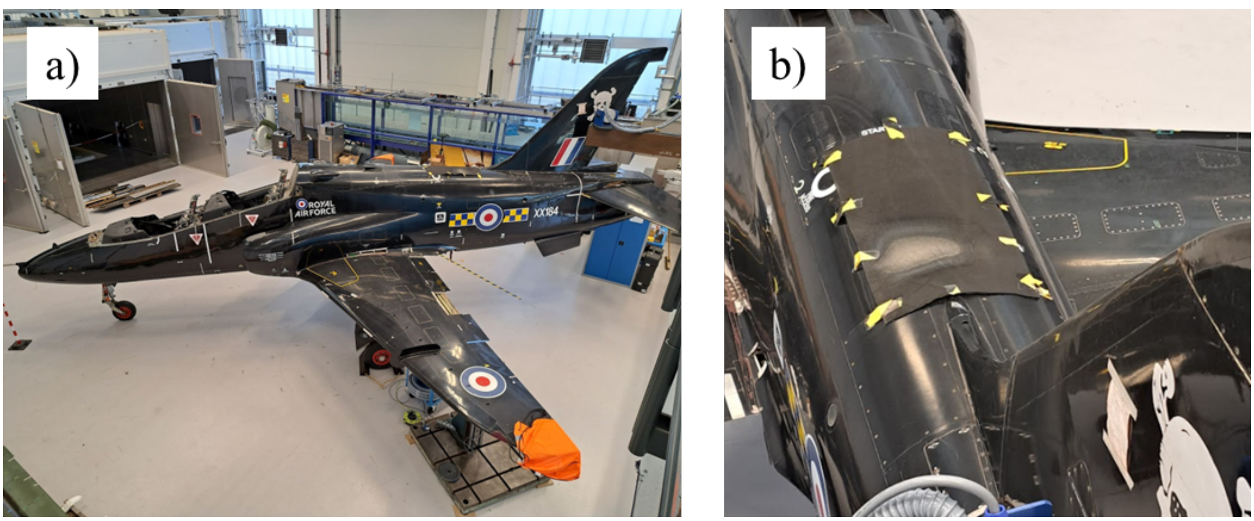

Fig. 1 a) The Hawk in situ and b) detail showing the missing upper fuselage, covered with fabric

The hawk was provided without documentation or a service history. All measurements and models included in this paper

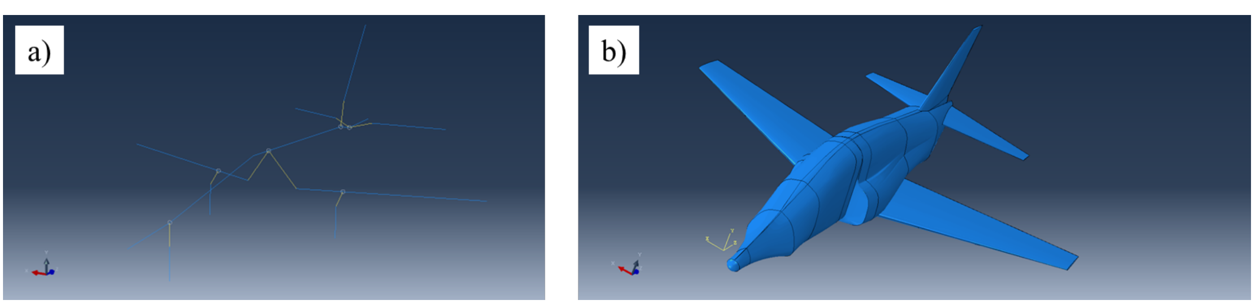

Fig. 2 a) The beam and b) shell CAD models of the hawk

are either from open sources or were collected during the project. To demonstrate digital twin capabilities, two computer aided design models were created of the Hawk in ABAQUS - a beam model and a rough shell model, shown in Fig. 2. The latter was also used as the basis of a 3D printed physical model of the Hawk, which aided discussion in the testing design phase.

Test Design

The Hawk was equipped with five shakers: one on each wing, one on each stabiliser, and one on the rudder. A total of 82 uniaxial accelerometers were applied to the fuselage, wings, stabilisers, rudder, and landing gear, as well as three triaxial accelerometers at the nose, cockpit and tail. Ultimately, the limitations were the number of channels available and the availability of the sensors themselves. An additional triaxial accelerometer was attached to the floor below the centre of the aircraft to detect ground vibration, while ambient temperature and sound were captured using a resistance temperature detector and microphone, respectively. Full detail about the test setup, including schematic diagrams of the locations of all sensors and shakers, can be found in Reference [3].

Fibre Bragg grating (FBG) was also applied along the lengths of the wings, stabilisers, and rudder. The reasoning for this was that FBG is lighter and less bulky than accelerometers. If damage could be successfully detected using the signals either the FBG and a reduced number of accelerometers (or even from the FBG alone) then this might present a more practical option for in-operation sensing.

The placement of the accelerometers was motivated by sensor placement optimisation [4]. Two separate methods were trialled: in this case effective independence [5] was found to be too computationally expensive and so the modal kinetic energy method [6] was used. The approximate nature of the underlying model (especially the fact it only represented the shell of the Hawk) meant that the applicability of the actual sensor locations calculated through this method was inherently limited. However, it did provide value by indicating a suitable distribution of the available sensors (i.e. how many should be placed on the wings compared to the tail). Expert knowledge, both from general experience with vibration testing and specific experience from the wing-only tests then guided the final placement of the sensors.

Data Collection, Storage, and Dissemination

The testing regime consisted of five phases, detailed in Ref. 3:

1. The Hawk in its normal, undamaged condition, using both white and pink noise inputs.
2. Single site pseudo-damage, using mass addition at one of fifteen locations and white noise input.
3. The Hawk in its undamaged condition, using odd random-phase multisine excitation.
4. Multi-site pseudo-damage with mass addition at up to three of the fifteen sites at once, using white noise input.
5. Damage test through the removal of panels on the port wing

A schematic of the test setup is shown in Fig. 3. The control equipment was operated using National Instruments (NI) Labview on the test PC. An analog output (A/O) card interfaced to the shaker amplifiers. The shakers themselves were then connected to the Hawk through stingers and load cells, with the load cells connecting back to the test PC through an analog input card. Analog input cards were also connected to the accelerometers, temperature and microphone, allowing the data from these sources to be captured through Labview synchronously with the output. A separate interrogator unit digitised the data from the FBGs and outputted it to the PC, where it was captured in a separate file to the NI data (although it was still acquired by the labview code). A synchronisation pulse was captured in all data files to ensure data alignment.

The multiple input, multiple output (MIMO) testing required significant effort at the start of each series of tests in order for each shaker to deliver the required signal. First, a breakpoint table was derived for each test sequence, describing the amplitude of the required excitation at each location. This was then resampled to create a breakpoint mask for each frequency line which, alongside a random phase, defined the sine wave for each frequency line, which combined to form the time domain signal. Each shaker was then excited, the response of the load cell compared to its expected value, and the gain mask updated accordingly. This process was then iterated until the target and measured values were within the specified tolerance of 4% and this final drive signal was used to gather the test data. New drive signals were generated using this method whenever any significant change might have occurred to the structure (i.e. at the start of each day and whenever masses or panels were added or removed) or whenever a different input signal was required.

Tests were logged using a Google sheet, in order to keep an up-to-date list across the multiple researchers involved, along with a Google form for redundancy and to capture any research notes. Data was then uploaded to the shared Google drive at the end of each day, as a backup and to allow for verification.

A total of 216 tests were conducted, creating in excess of 500 GB of raw data. In order to ease the practical challenges associated with working with a dataset of this scale, the dataset is packaged using the *hierarchical data format* (.hdf5), which allows the data and metadata to be stored concurrently. In addition, a simple interface was developed in Python to aid users. This API first checks for local versions of the requested files and then, if not available, accesses the data from a central repository. This means that users only download the exact data required and avoids data erosion through repeated sharing of the data.

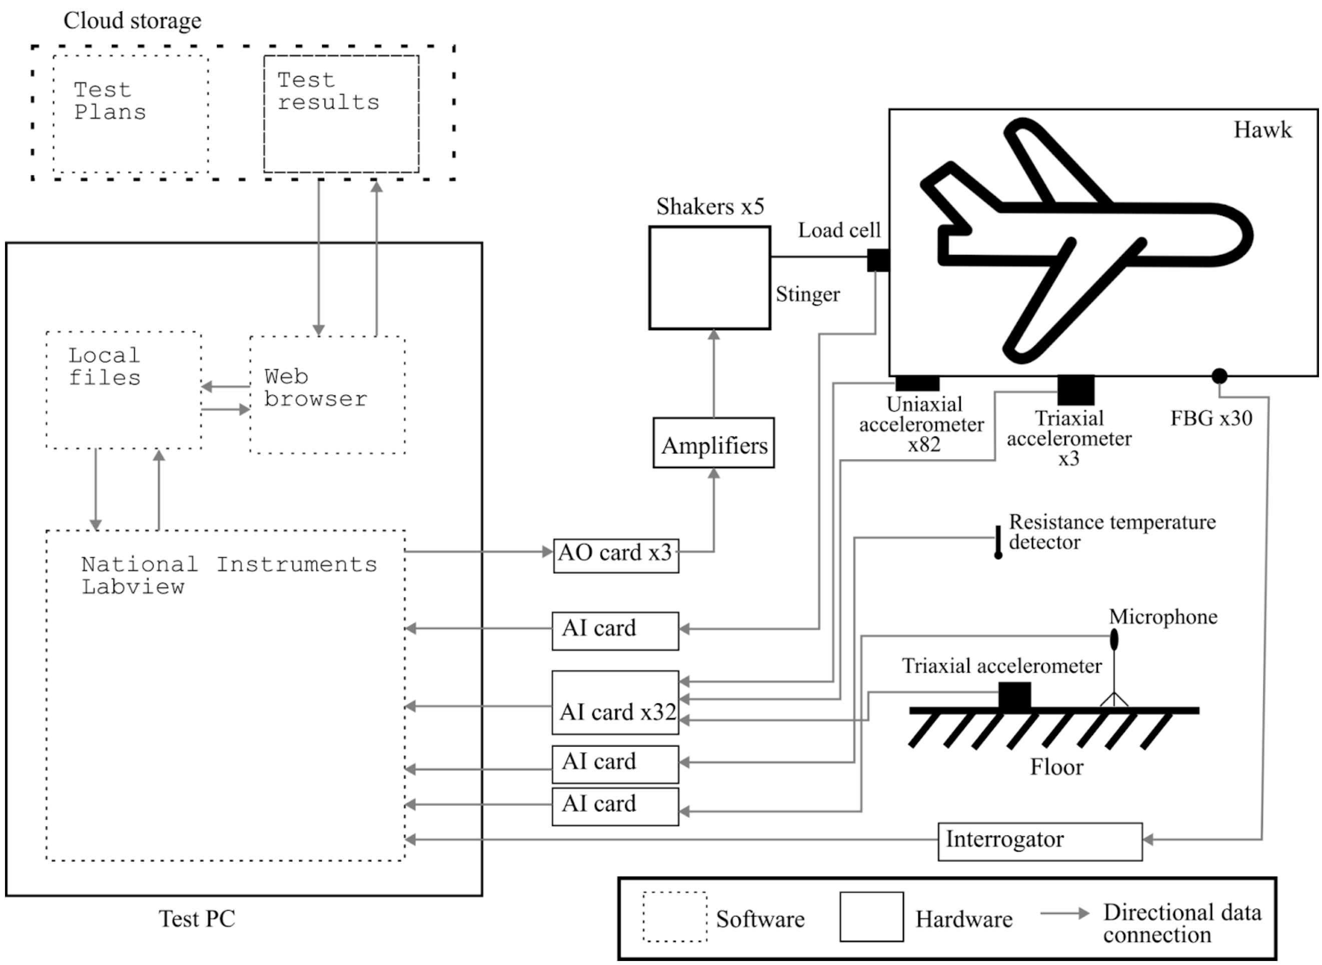

Fig. 3 A schematic of the test setup, showing both software and hardware components

Digital Twin Online Platform

An implementation of the Digital Twin Operational Platform (DTOP) [7] was used to demonstrate some digital twin capabilities. DTOP is an open source, browser-based digital twin architecture developed using Python and Flask. The DTOP platform enables the creation of a web-page or web-site, containing data and models of the physical twin, to act as a single source of truth. Currently these web-pages are created and deployed locally. In this instance, the DTOP interfaces with the data asynchronously, with real-time data linkage being a topic for future work.

The specific implementation of the DTOP created for the Hawk is nicknamed DTHive and is shown in Fig. 4. The DTHive webpages offer the user two capabilities. The first is access to the data from the experimental tests. The second capability is interfacing with the ABAQUS finite element analysis software to investigate and update the CAD models. This

Fig. 4 The landing page of DTHive

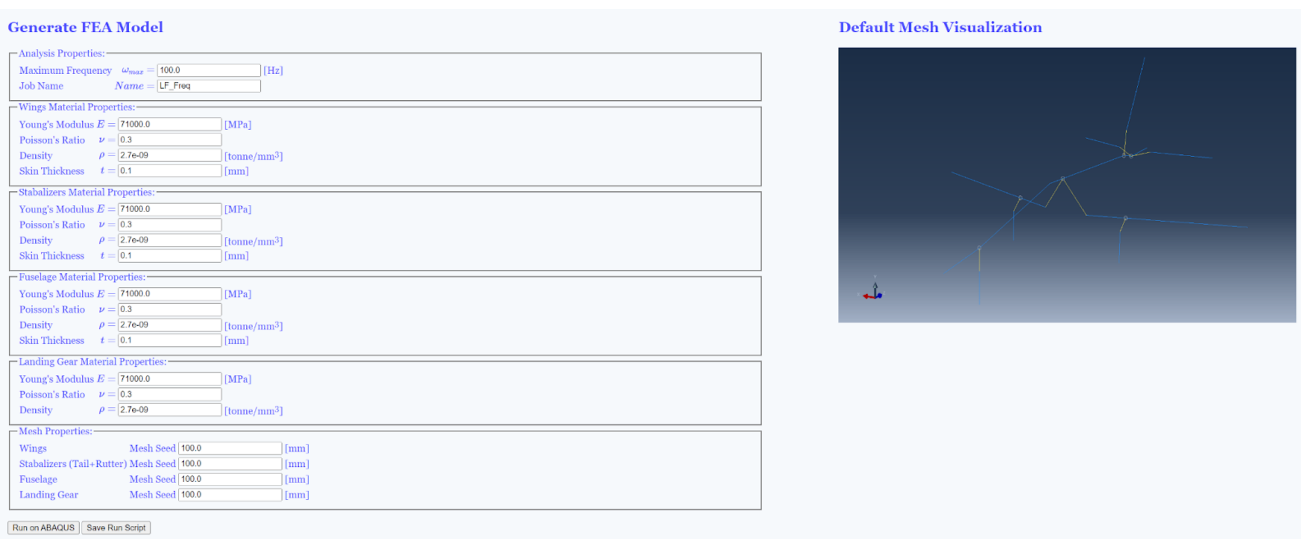

Fig. 5 The beam model generation page of DTHive

requires the user to have a local, licenced copy of ABAQUS. Running the script through the interface in Fig. 5 allows the user to open either the beam or the shell models from the DTOP, edit the material properties and load the resulting file into ABAQUS.

Data

Figure 6 shows an indicative sample of the data collected using the accelerometers and FBG sensors, respectively. An example of some pseudo damage data is shown in Fig. 7, which shows the response of an accelerometer on the starboard stabiliser to white noise excitation with the addition of the largest mass (M1, 245.3g) and smallest mass (M7, 64.4g), as well as without mass addition. Further detail can be found in Reference [3], which also includes a number of motivating research

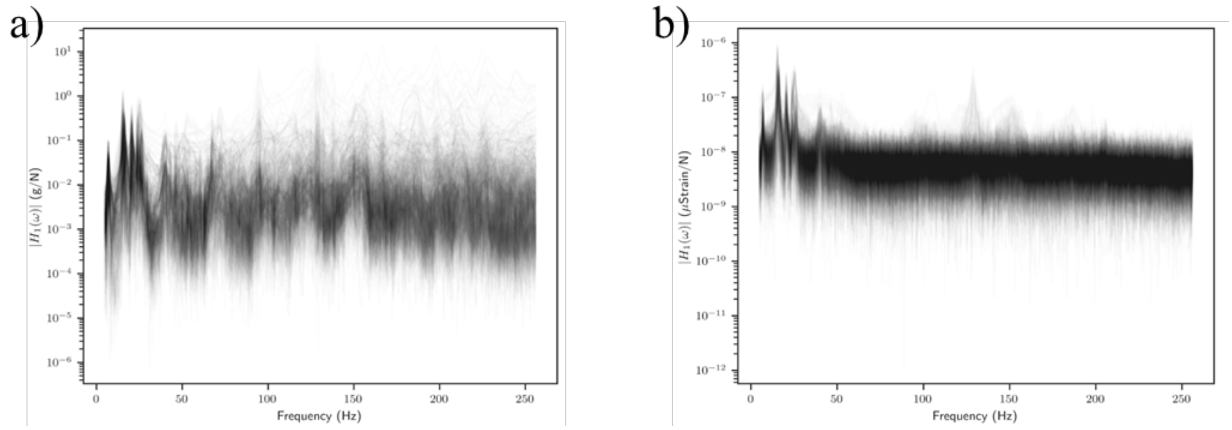

Fig. 6 Superimposition of the frequency response function magnitudes for white noise healthy state tests from a) all accelerometers and b) all FBGs

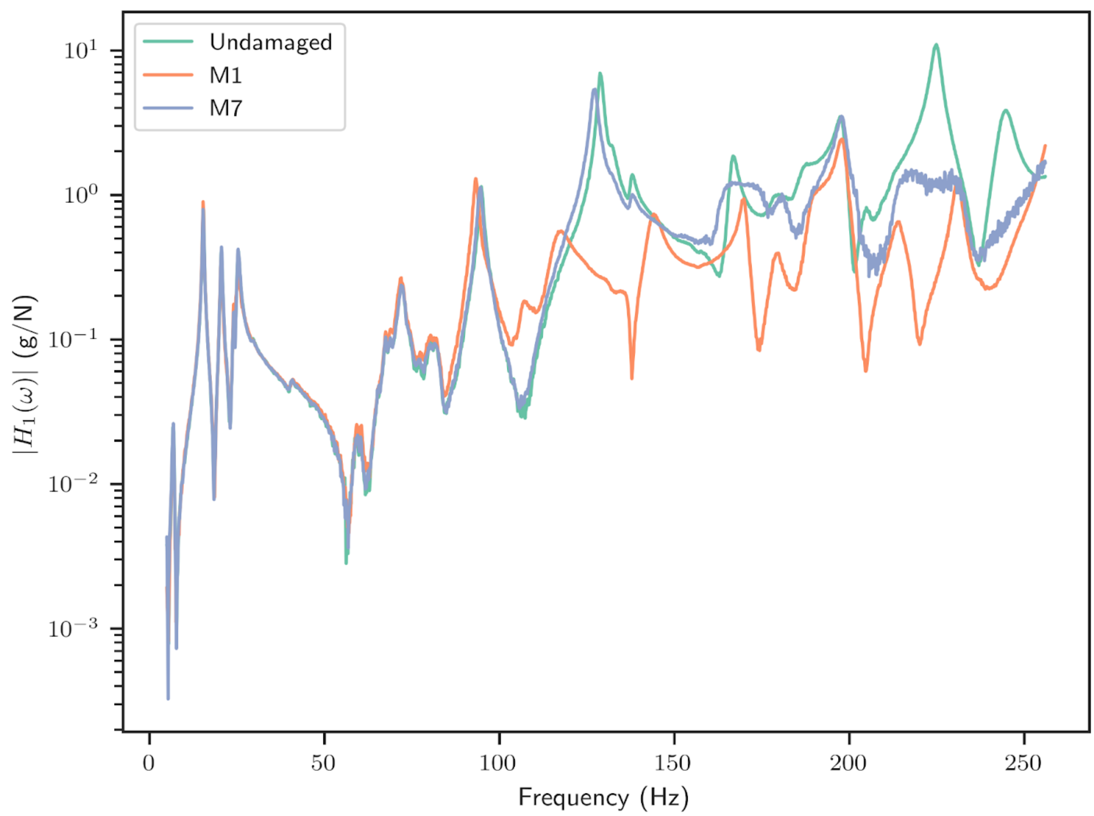

Fig. 7 Frequency response function magnitude of a single accelerometer on the starboard stabiliser in response to white noise input in an undamaged state and in the presence of two different masses

challenges. Of particular relevance to the discussion here are Challenge 11: Detecting pseudo-damage using subsets of the available data, Challenge 13: Detecting pseudo damage with exclusively the FBG data, and Challenge 14: Establishing the extent to which higher quality data is useful in the training phase, when only lower quality data is available during operation.

Discussion

The collection of the Hawk dataset highlights a number of key issues relating to the creation of digital twins. These largely relate to the nature of the purpose of the digital twin, and the identities of the users and data owner. While many aspects of digital twin design are widely applicable, as we move towards more industrially applicable demonstrators, the specifics of each of these aspects become more relevant and we are required to make assumptions about each. These assumptions radically affect the specifics across the design, including the sensor selection, data pipelines, modelling, and data visualisation.

For land-based structures, such as bridges and wind turbines, accelerometers and strain gauges are standard tools for monitoring structural dynamics. However, for aircraft, weight reduction and aerodynamics are key considerations. Attaching over 100 sensors to the external structure of an operational aircraft without significantly degrading performance would be at minimum a significant design challenge. Equally, attaching such a large quantity of sensors, along with their associated cabling, requires multiple days to achieve, which makes gathering as dense a dataset between flight hours impractical. These concerns motivated the inclusion of the fibre bragg grating in these tests, as these lower profile, lighter sensors might be more practical. However, resilient damage detection methods need to be shown using this FBG data. Similarly, if data from subsets of the accelerometers can be used for damage detection, then this would present a significant step towards the use of sparse data. The specific requirements of the digital twin matter here: if in-operation damage detection is required, then integration of light weight sensors is key; if instead the goal is to support the scheduling of on the ground maintenance, then effective sensor placement optimisation or investigation of other methods, such as the use of robotic laser vibrometry must be prioritised.

In both cases, the data collection and modelling are best supported by a high level of confidence in the knowledge of the structure in its undamaged state. The limitations of conducting sensor placement optimisation using the approximate models of the Hawk have already been discussed. High fidelity models also have value in other areas, such as interrogating the causes of nonlinearities in the underlying data. However, the generation of such models is not trivial, even with recent advances in 3D scanning, as is gaining a bill of materials for a sensitive legacy asset like the Hawk. If the involvement of the original equipment manufacturers in the creation of the digital twin can be assumed, then these issues become trivial. Otherwise, however, decisions will need to be made regarding the amount of effort to be spent on generating high quality computer models, and so the benefits need to be understood and well evidenced.

As is often the case with prototyping and experimental work, the data pipeline used here was low in automation, using immensely more human labour than would be practical in a real-world context. This includes both time in the set-up stage and during the testing itself. The requirement to check the attachment of the sensors manually, followed by a lengthy generation of the drive signals before each series of tests required a level of human involvement that precluded the automation desirable from a digital twin. Inclusion of a SCADAS style sensor integrity testing of accelerometers, as an example, could massively decrease the effort involved. Including data verification earlier in the pipeline would also be of major benefit, allowing for errors to be caught and tests to be re-run more quickly. However, given the nature of collecting a dataset that, by design, includes damage that is hard to distinguish, only so much verification can be automated. The number of edge cases mean that expert knowledge to distinguish when data "looks a bit off" is still required. These issues are included in Reference [3] as Challenge 2: Challenge 2. Aiding (or reducing the reliance on) operator interpretation when collecting large datasets.

Many improvements were made in the testing protocol, building on the lessons learned from the single wing tests. One example is the inclusion of the microphone and resistance temperature detector. Having a comparable record of ambient testing conditions massively reduced the number of subjective, hard to quantify comments included alongside the data, as testers were no longer required to judge what levels of additional ambient sound were notable. As digital twin tools develop, this kind of ability for continuous improvement in workflows should be kept in mind, with tools designed to encourage them rather than lock in suboptimal workflows.

A final comment concerns data visualisation. In this project, relatively conventional data visualisation techniques have been used, as the dataset is designed to be used by subject matter experts. However, a major promise of digital twins is the reduction of silos and the greater dissemination of data across organisations. The ways in which structural dynamics data might be best conveyed to non-expert audiences will need to be studied, especially where it relates to uncertainty quantification and conveying the limits of what the data *can* show (i.e. avoiding overconfidence). As ever, the users and purposes need to be considered in all this - a ground technician crew will have very different understanding and requirements to a pilot in flight.

Conclusion

This paper has described the collection of the Hawk structural dataset from the perspective of digital twins. The structure, test design, test equipment and data collection were described in detail. A version of the digital twin online platform software representing the Hawk was also described. Finally, some discussion was given to the limitations of this testing as it relates to practical implementations of digital twins. This discussion focused on the ways in which digital twin design is necessarily motivated by the purpose of the twin, the identity of the user, and the identity of the data owner and covered aspects of sensor delection, the difficulties of creating digital twins of legacy structures, data pipeline design, test automation, the need for continuous improvement, and dat visualisation.

Data Availability

All data collected is made freely available at: orda.shef.ac.uk/articles/dataset/BAE_T1A_Hawk_Full_Structure_Modal_Test/24948549. In addition, the authors make a python interface for interacting with the data available at: https://github.com/MDCHAMP/hawk-data.

Acknowledgments This work was funded by the Alan Turing Institute Research and Innovation Cluster for Digital Twins and EPSRC grant EP/R006768/1. The authors gratefully acknowledge this support. The authors would also like to acknowledge the many individuals who gave their time and effort to collect this data.

References

1. D. J. Wagg, K. Worden, R. J. Barthorpe, and P. Gardner, "Digital Twins: State-of-the-Art and Future Directions for Modeling and Simulation in Engineering Dynamics Applications," ASCE-ASME journal of risk and uncertainty in engineering systems, Part B. Mechanical engineering, vol. 6, no. 3, 2020, doi: 10.1115/1.4046739.
2. M. Haywood-Alexander et al., "Full-scale modal testing of a Hawk T1A aircraft for benchmarking vibration-based methods," Journal of sound and vibration, vol. 576, p. 118295, 2024, doi: 10.1016/j.jsv.2024.118295.
3. J. Wilson, M. D. Champneys, M. Tipuric, R. Mills, D. J. Wagg, and T. J. Rogers, "Multiple-input, multiple-output modal testing of a Hawk T1A aircraft: A new full-scale dataset for structural health monitoring," 2024, doi: 10.48550/arxiv.2406.04943.
4. S. Hassani and U. Dackermann, "A Systematic Review of Optimization Algorithms for Structural Health Monitoring and Optimal Sensor Placement," Sensors (Basel, Switzerland), vol. 23, no. 6, p. 3293, 2023, doi: 10.3390/s23063293.
5. D. C. Kammer, "Sensor placement for on-orbit modal identification and correlation of large space structures," in Proceedings of the American Control Conference, 1990, pp. 2984–2990. doi: 10.23919/acc.1990.4791265.
6. M. Meo and G. Zumpano, "On the optimal sensor placement techniques for a bridge structure," Engineering structures, vol. 27, no. 10, pp. 1488–1497, 2005, doi: 10.1016/j.engstruct.2005.03.015.
7. M. S. Bonney et al., "Development of a digital twin operational platform using Python Flask," Data-Centric Engineering (Online), vol. 3, no. 1, 2022, doi: 10.1017/dce.2022.1.