Tools and Techniques for the Maintenance and Support of Digital Musical Instruments

Albert-Ngabo Niyonsenga
IDMIL, CIRMMT
McGill University
Montreal, QC
albert-ngabo.niyonsenga@mail.mcgill.ca

Marcelo Wanderley
IDMIL, CIRMMT
McGill University
Montreal, QC
marcelo.wanderley@mcgill.ca

ABSTRACT

There are multiple barriers to the long term use of digital musical instruments. Among several issues related to instrument accessibility, many DMIs remain as prototypes in research labs never becoming a robust and stable instrument. Technical support is an important part of the long term use of a DMI. Though all musical instruments can eventually break, managing how they are going to be fixed and built within a research organisation can help with the continued usage of the instrument. We apply reliability analysis techniques to estimate the reliability, availability and maintainability characteristics of the T-Stick. Using these characteristics we estimate the amount of spare parts needed to maintain a 99% availability target for the T-Stick. This analysis provides insights on expected maintenance time, costs, and personnel needed when supporting and building DMIs.

Author Keywords

T-Stick, DMI, reliability

CCS Concepts

•Applied computing → Sound and music computing; •Hardware → Hardware reliability;

1. INTRODUCTION

All musical instruments require maintenance and repair at some point in their lifetime. Acoustic instruments typically have trained instrument repair technicians/luthiers who are able to assist musicians with keeping their instruments in good shape. Digital Musical Instruments (DMIs), mostly those designed in research organisations, do not typically have this. Such instruments face several challenges hindering their long term use. Though a large number of DMIs are presented at the New Interfaces for Musical Expression (NIME) conference every year, only few of them remain in use due to a large variety of issues.

Although for many devices presented at NIME it may be appropriate for the instrument to remain ephemeral [5]and fade into obscurity after its initial artistic/research contribution has been documented, in many cases there is a need for stable and robust instruments which can be reliably used for longer periods of time. Indeed, "most of the instruments have difficulties establishing themselves" [11]. Similarly, "most NIMEs are viewed as exploratory tools created by and for performers, and that they are constantly in development and almost in no occasions in a finite state" [12].

Supporting a DMI requires one to consider many variables. Who has built the instruments and where they were built? Who is going to fix them? How much time will it take to fix these instruments? How many spare parts/instruments are needed to allow performances/projects to go forward while repair is underway?

Though it may be argued the "interface construction proceeds as more art than science, and possibly this is the only way that it can be done" [3], if longevity is the goal of an instrument's design, then a balance between art and engineering is required to achieve reproducible, robust and reliable instruments [10].

Transitioning from a laboratory prototype to a stable instrument is neither easy or straightforward. Although a prototype instrument with reliability issues may be tolerated a stable instrument would be expected to perform well under various performance conditions. Evaluating this is not straightforward and can be hard to quantify. Evaluation methodologies from other fields such as Human Computer Interaction [15] have been applied to the use of the design and evaluation of DMIs. For reliability we can look look towards reliability engineering and how this field approaches reliability.

In this article we reflect on our experience with building and maintaining the T-Stick and the lessons that were learnt in the process. We then apply *reliability*, *availability* and *maintainability* (RAM) analysis and *spare parts optimisation* to evaluate the reliability and maintainability the T-Stick [9]. The T-Stick is used as it is an instrument that is transitioning between prototype and a stable product. A RAM analysis on the T-Stick provides a formal means to quantify the current reliability of the T-Stick and identify issues with the current design and manufacturing process. These techniques are used to estimate the amount of spare parts and T-Sticks necessary to maintain a 99% availability target for this DMI. This allows one to better estimate the maintenance needs and cost of their instrument ahead of time.

2. T-STICKS

The T-Stick is a gestural controller designed by Joseph Malloch in collaboration with Andrew Stewart and Marcelo Wanderley [9]. It is a family of DMIs with multiple dif-

NIME'23, 31 May–2 June, 2023, Mexico City, Mexico.

ferent versions differing based on their length and the total amount of sensors. These include the Tenor, Alto, Soprano and Sopranino T-Stick. As of January 2023 there are currently 1 Tenor T-Stick, 6 soprano T-Sticks, and 16 sopranino T-Sticks at the Input Devices and Music Interaction Lab (IDMIL).

They are composed of an inertial measurement unit (IMU) and multiple pressure, piezoelectric and capacitive sensors. The most recent version of the T-Stick uses the Trill craft board for the capacitive sensors. Older versions as shown in fig.2 use a custom capsense board [1]. Bigger versions of the T-Stick differ from the sopranino mostly by requiring more capacitive sensors and force sensitive resistors. The newer versions of the T-Stick built after October 2021 no longer use the piezoelectric sensors.

These sensors extract gestural information related to the T-Stick. This includes gestures such as squeezing, twisting and brushing as well as the raw data from the sensors. Furthermore the instrument also sends its orientation and can detect jabs and shaking[2]. The T-Stick is typically used with Libmapper or OSC in order to connect it to different DMIs and sound synthesizers. Libmapper is a tool that facilitates conections between DMIs on a network [8].

3. OUR EXPERIENCE BUILDING AND REPAIRING DMIS

In the Fall of 2021, a new batch of T-Sticks were built as part of a class assignment. Due to previous components of the T-Stick no longer being in stock, the hardware had to be quickly redesigned to maintain the same functionality. 9 T-Sticks were built as part of this assignment, 1 Tenor T-Stick, 1 alto T-Stick, 2 Soprano T-Sticks and 7 Sopranino T-Sticks. Most of the builders had little to no experience with electronics. At the end of the 4 weeks only two T-Sticks were functional.

It took 3 months working approximately 5-10 hours a week to get 7 of the T-Sticks somewhat operational. The T-Sticks were able to turn on and send signals via OSC/Libmapper but still experienced failures when shaken or tapped vigorously. Some manufacturing defects of these T-Sticks were so severe that it would often require a near complete rebuild in order to ensure it functioned properly. This is both time consuming and unsustainable in the long term if T-Sticks are to continue to be built in-house.

The failures encountered included, but were not limited to, sensors or TinyPico boards not responding, the battery not charging or holding charge for a short amount of time, and the T-Stick rebooting upon being shaken. The process that was undertaken while fixing the instrument involved first identifying the fault, attempting to fix the fault and then testing that the T-Stick functioned. This last step proved to be the most difficult as there was not an established testing methodology for the T-Stick. There were a series of quick test that checked some functionality of the T-Stick but not others. It was easy to test if sensors were sending data but testing if that data was accurate proved challenging given the tools and time available.

Although this is an extreme example this case highlights some of the dangers of in-house manufacturing and maintenance. Unlike a professional assembly house there are fewer resources to catch and minimise manufacturing defects. This can lead to more variance between the build qualities of each instrument.

That is not to say there are not good reasons to consider building the instrument in-house. By taking responsibility for building the instrument the organization gains valuable technical skills. It can be used as a teaching tool for students and new researchers to gain practical experience with building DMIs. Over time supporting the instrument becomes easier as common solutions to issues become known and documented. This knowledge can be used to improve the design of the instrument in future iterations or used to improve the design of new instruments.

Furthermore, using an external partner is not risk free. Even if the manufacturer is experienced in manufacturing electronics it does not mean that they will be able to manufacture T-Sticks at a higher quality than students or researchers. In addition there are always supply chain risks that can lead to significant delays in receiving the instrument.

4. MAINTENANCE TARGETS

Regardless of whether the instrument is maintained in house or by an external partner, one must set maintenance targets/goals. Part of the difficulty with maintaining and fixing the T-Sticks were not only that some were hastily built, but that the design standard of the T-Stick was ignored. Ignoring the design standard for the T-Stick while building them means that several T-Sticks may have been able to turn on but did not perform well.

Although it may seem that a standard is overkill, what a working DMI looks like can vary quite substantially. For example with the T-Stick, at what point does the battery need to be replaced. Is it after it cannot last 8 hours under a single charge, 2 hours, 30 minutes, etc. How hard does the T-Stick need to be shaken? If a light shake is fine but a vigorous shake is not, is that still considered a "working" T-Stick? Does the signal accuracy and precision matter in terms of a working instrument? If the sensor drifts with time, how much drift is seen as acceptable before it is no longer considered a working sensor?

Furthermore, although the operating conditions of an instrument may seem clear at first, there are a lot of details that need to be considered. This includes how long it needs to operate in a single session, if it is designed to work in an indoor or outdoor venue, how many days of the year is the instrument going to be used. Artist and performances can have varying rehearsal schedules, practice times, and practice frequencies.

Ultimately the organisation has to understand these varying conditions and update their analysis accordingly. Assuming the worst case scenario will mean taking on a greater cost in terms of spare parts and more time and people will need to be allocated to maintain the instrument. Alternatively, a more thorough analysis may be done to optimise the maintenance schedule to better reflect the conditions. Understanding this target is important as maintenance and support of instruments cost your organisation both time and money. Being able to estimate the time and money ahead of time will allow the organisation to better plan how the maintenance of their instrument will happen.

5. RELIABILITY ANALYSIS

To understand the maintenance and support costs in terms of both dollar cost and time cost we must first understand the reliability and maintainability of our instrument. By understanding reliability and maintainability, we can estimate how much it will cost to maintain a target availability.

[1]The full bill of materials to build a Sopranino T-Stick can be found at `https://github.com/IDMIL/T-Stick`

[2]A full list of the data the instrument sends can be found on at `https://github.com/IDMIL/T-Stick`

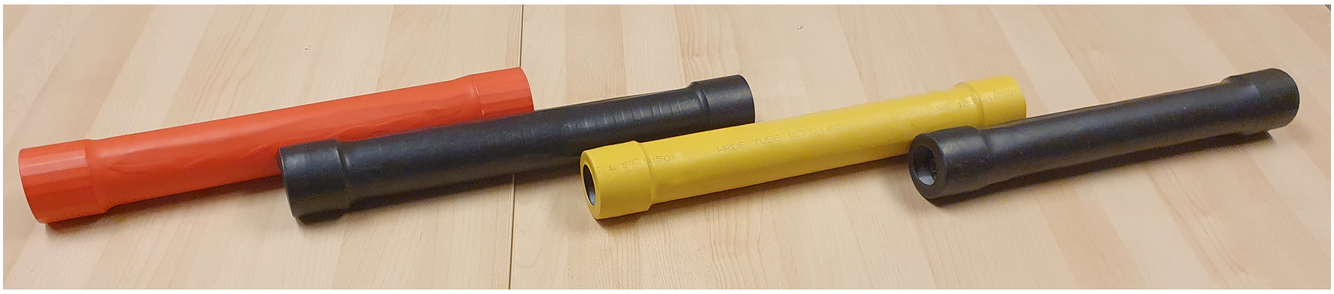

Figure 1: Four Sopranino T-Sticks

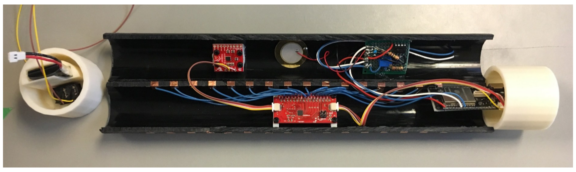

Figure 2: Inside a Sopranino T-Stick, picture from Github. URL: `https://github.com/IDMIL/T-Stick/blob/master/Docs/T-Stick_2GW_building_instructions.md`. Accessed 18 Oct. 2021

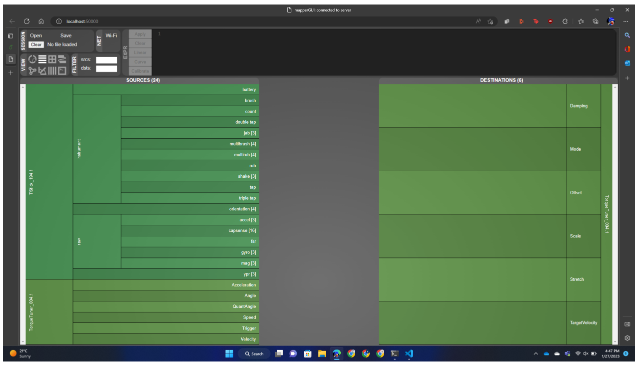

Figure 3: Example Webmapper interface, with T-Stick output signals on the left in dark green

This can be done by conducting a Reliability, Availability and Maintainability (RAM) analysis on the instrument.

RAM Analysis consists of estimating availability of the device by using the reliability and maintainability of the individual components. With this information one can estimate how much support the device will require to operate and estimate how much it will cost [13]. In order to do this one must be able to estimate the reliability and maintainability of the device.

Reliability consists of three parts. It is a measure of a device's ability to perform a specific function under defined environmental and operational conditions over a defined period of time [1]. Reliability is measured by probability of the system to not fail within a given time. Reliability is assumed to follow a "bathtub curve" with the failure rate being high at the early stages of the devices life cycle due to undetected manufacture defects and later stages of the life cyle of the device due to wearing out of components and constant during the middle of the life cycle.

Assuming a constant failure rate we can compute the *mean time between failure* (MTBF)/*mean time to failure* (MTTF) by taking the reciprocal of the failure rate. Note that this estimate of mean time between failure only considers random failures.

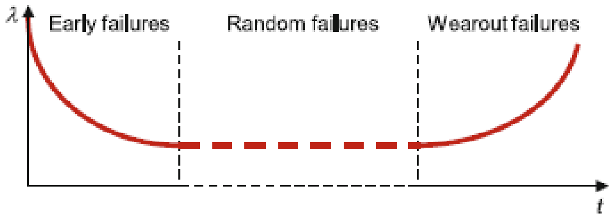

Figure 4: Bathtub Curve [7]

$$MTBF = \frac{1}{\lambda} \tag{1}$$

where λ is the failure rate.

For a constant failure rate λ we can compute the reliability as an exponential distribution.

$$R(t) = e^{-\lambda t} \tag{2}$$

Any systemic failures are ignored by this analysis. Although this assumption is generally good, it is hard to prove that for any device a majority of the failures are random [7]. This means that this estimate for reliability can be an underestimation of failure in the system.

Maintainability is a measure of how easy the device is to fix. We measure this using *mean time to repair* (MTTR). Furthermore, maintainability can be extended to consider lead time of components. *Availability* is a combination of both reliability and maintainability. It is a measure of average up time of the system [7]. If only the downtime related to fixing the instrument is considered then that is called *inherent availability*.

$$IA = \frac{MTBF}{MTBF + MTTR} \tag{3}$$

where IA is inherent availability.

Achieved availability considers the impact of the lead time of components and *operational availability* considers all sources of down time including administrative sources [6].

Alternatively, one can consider an alternative metric dispatch availability. Dispatch reliability is the measure of how often one can dispatch an instrument when it is demanded. This metric is often used in aerospace, where we consider how often we can dispatch a plane and the average delay involved [2]. For an instrument one consider how often we can dispatch instruments to performers/experimenters when they need them.

$$DR = 1 - \frac{\text{delay} + \text{cancellations}}{\text{total requested dispatches}} \tag{4}$$

where DR is dispatch reliability.

Dispatch reliability offers us a different perspective on the impact of reliability of the device on the maintenance costs to the organisation. However, it should be noted that unlike airplanes one rarely has to give out instruments on a set schedule for a set amount of time. Outside of performances and experiments instruments can be assumed to not be in use.

Reliability block diagrams (RBDs) are used to show the relationship between components. These diagrams are a visual representation of failure relationships between components. Any component in series means that a failure in that component causes a failure in the system. Components in parallel represent redundancy in the system, allowing some of the components to fail without causing an entire system failure. For simple series systems we can compute the RAM characteristics of the system by multiplying the reliability of each component together. We can also compute the availability of the device in the same manner. To compute the maintainability of the system one must rearrange Eq. 3 to solve for mean time to repair.

Consider a simple RBD shown in fig. 5. It is made up of three components A , B and C with MTBF of 100 hours, 50 hours and 25 hours respectively. To compute the RAM characteristics of the device we first compute the failure rate of each component by rearranging Eq. 1. This gives us a failure rate of 1%, 2%, and 4% for each component. To compute the reliability of the entire system we have to multiply the reliability of each component together. This is the equivalent of adding all their failure rates together. Doing this, we get a failure rate of 7% and a MTBF of 14.57 hours.

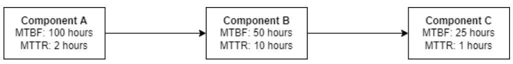

Figure 5: RBD of a simple device with no redundancy

To compute the availability of each component using Eq. 3 we first compute the availability of each individual component using Eq. 3. This gives us an availability of 98.0%, 83.3%, and 96.2% for each component. To compute the total availability we multiply these availabilities together which gets us an availability of 78.6%. To compute the system MTTR we re-arrange Eq. 3 to solve for MTTR and obtain a system MTTR of 3.97 hours.

From this simple analysis one can see that the availability and reliability of the system is never going to be better than the reliability and availability of the individual components. We note that in this case component B has the lowest availability and is contributing to this low overall availability.

What if we add some redundancy to the system? We can consider the system shown in fig. 6 where there are now three copies of component B in parallel. In order for the system to fail all three copies must fail.

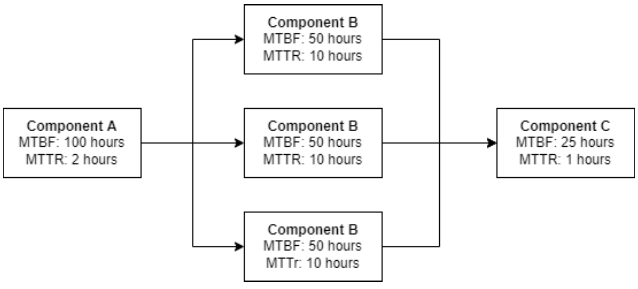

Figure 6: RBD of simple device with some redundancy

To find the availability and reliability of this parallel system we first compute the RAM characteristics of the three component Bs. To do so we need to compute the probability that they will all fail at the same time. The more generic version of this problem is given by Eq. 5.

$$R_s\left(t\right) = \sum_{k=m}^{n} \binom{n}{k} \left(r_c\left(t\right)\right)^k \ \left(1 - r_c(t)\right)^{n-k} \tag{5}$$

where R_s is the reliability of the redundant system, r_c is the reliability of a component, n is the total amount of components, and m is the minimum amount of working components.

Assuming $t = 1$, we get a failure rate of about 0.0008% or an MTBF of 128800 hours. One can see here that in any system with redundancy, the failure rate will always be lower than the smallest individual failure rate. We can estimate the MTTR of the three components by multiplying the MTTR by three. Now that we have the MTBF and MTTR of the three component Bs we can proceed as we did for the series system and compute the reliability, availability and maintainability of the total system. This gives us a MTBF of 33.54 hours, a MTTR of 7.52 hours and an availability of 81.68%. The device with redundancy is twice as reliable failing on average every 33.54 hours in comparison to every 14.57 hours which leads to a higher availability. However, it has a higher mean time to repair due to the additional components.

5.1 RAM Analysis of a T-Stick

For this analysis we are concerned with hardware failures that will require the instrument to be fixed in an electronic lab. Furthermore, we consider that the T-Stick will be used in 8 hour sessions. We consider a T-Stick to be available when there are at least three functioning T-Sticks. We do this because performances can involve multiple T-Sticks therefore understanding the maintenance load of maintaning multiple T-Sticks is crucial. We aim for an availability target of 99%.

In order to conduct a RAM analysis on the most recent version of the Sopranino T-Stick we must consider the function of a T-Stick and the environmental conditions the T-Stick will operate in. As a gestural controller, the T-Stick's main function is to send sensor data to a separate system for the purpose of synthesising sound signals. In addition, the T-Stick sends both raw sensor data and high-level gestural data such as jabs, rubs and brushes. We will only consider a T-Stick failure as a fault that stops the T-Stick from sending signals over WiFi.

The T-Stick can be broken into three subsystems; a control system, sensor system and power system. The control system includes a TinyPico development board along with solder joint connections for the I2C wires, the push button, the force sensitive resistor (FSR), and battery. The sensor system is made up of all the sensors in the T-Stick. This includes a Trill board, LSM9DS1 IMU, a push button and an FSR. The power system is made up of the battery holder and the 18650 lithium ion battery. Each system consists of solder joints connecting components together. The Fall 2021 T-Sticks have approximately 51 manual solder joints. Table 1 shows the reliability parameters of each component. These parameters are estimates based on the MIL-HDBK-217F handbook [4] our experience with the T-Stick and with each of these components[3].

[3]A more formal analysis should use more detailed figures from standards such as Siemens SN 29500 [14] or from test on the instrument itself.

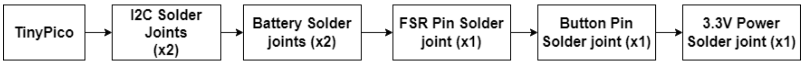

(a) Reliability Block Diagram of the Control System of a Sopranino T-Stick

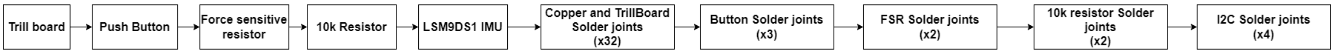

(b) Reliability Block Diagram of the Sensor System of a Sopranino T-Stick

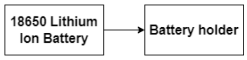

(c) Reliability Block Diagram of the Power System of a Sopranino T-Stick

Figure 7: Reliability Block diagrams of the subsystems of a Sopranino T-Sticks

Table 1: RAM Characteristics for Individual Components

Component	MTBF (hours)	MTTR (hours)	Availability (%)
ESP32 Module	50000	3.5	99.9993%
Capsense module	50000	3.5	99.9993%
Battery	4000	1	99.9750%
LSM9DS1 IMU	50000	3.5	99.9993%
10k Resistor	9000000	1.25	99.9999%
Battery holder	20000	0.75	99.9995%
Push button	20000	1	99.9995%
FSR	20000	1	99.9995%
Solder joint	5000	1	99.9800%

The current design of the T-Stick has no redundancy which simplifies the reliability block diagram of each subsystem. Every component is connected in series with every other component. This means that any component failure will also cause a system failure. A partial RBD is shown in fig. 7. To calculate the RAM characteristics of each subsystem we use the same principles that we used to calculate the RAM characteristics of a simple series system. In the case of the T-Stick we extend that analysis for 59 components (51 solder joints, 8 other parts).

Table 2: RAM Characteristics of a Sopranino T-Stick

T-Stick RAM Characteristics	
MTBF (hours)	94
MTTR (hours)	1.01
Availability (%)	98.85%

The RAM characteristics of the Sopranino T-Stick can be found in table 2. Given that the inherent availability of a single T-Stick is 98.85% we can compute the availability of 3 T-Sticks by multiplying the availability of the 3 T-Sticks together. Therefore the inherent availability of 3 T-Sticks without any spare T-Sticks would be 96.85%. This is lower than the desired target of 99% and therefore implies that we need spare T-Sticks to reach our target. Spare T-Sticks allow us to introduce some redundancy. We can then compute the availability of the T-Sticks using Eq. 5.

Using this equation we can compute how the availability changes given the number of spares that we have. These results are in table 3.

Table 3: Inherent Availability of the T-Stick with spares

Spare T-Sticks	MTBF (hours)	Availability (%)
0	94	96.85
1	1509	99.80
2	85748	99.99

One can see that a single T-Stick has a mean time to failure of 94 hours. Assuming a T-Stick is used in 8 hour sessions for performances, this means we expect a single T-Stick to last about 12 performances before it fails. Alternatively, if we used 12 T-Sticks in a single performance we would expect 1 to fail. For the new batch of T-Sticks this estimate is rather optimistic. It is unclear whether the issue is just poor manufacturing or whether there are design flaws that cause early failures but the failure rate of newer T-Sticks is much higher than this figure would suggest. The most common failure mode of the newer T-Stick are I2C errors caused by solder joint failures. In addition other common faults include wire shorts which were due to poor manufacturing and improperly secured batteries.

Given the lack of space for proper wire management and the lack of experience of the builders of the T-Stick and that the solder joints in the T-Stick are put under a considerable amount of stress during the final process of assembly meant that many solder joints either failed immediately or were already significantly likely to fail given a moderate amount of stress such as shaking the T-Stick.

This is one of weaknesses of RAM analyses. They do not capture “systemic failures”, i.e. failures due to design choices or manufacturing defects are not captured by this analysis. However, by doing this analysis it can help point to these systemic failures. If the instrument is failing at a higher rate than is anticipated, it may point to design or manufacturing failures. An analysis can be conducted on why it is failing and the parts of your design that are causing the instrument to fail early.

Furthermore, this analysis does not consider the impact of lead time, assuming that we always have the spare parts required and a technician able to fix the instrument on hand. This is why were able to achieve such high availability given so little T-Sticks. Realistically, it is rare that a technician is immediately available to fix a T-Stick as soon as it arrives in the lab and there is travel time for the T-Stick to arrive to the lab from the performer or researcher that was using it.

5.2 Spare parts optimisation

The analysis above can be extended by simulating the impact of lead time and worker availability on the operational availability and therefore giving a more accurate estimate on how many spare parts are needed. Estimating how many spare components are needed is not a trivial task. Too few spare components increases the risk of the instrument not being available when it is needed, which can impact performers and researchers. Alternatively, having too many spare parts is costly.

One can guess the amount of spare parts needed based on their experience with the instrument and the components within it and how long it takes to get spare parts. Generally more spares are required for components that have a longer

Table 4: Lead Time of components

Component	Lead time (days)
ESP32 Module	7
Capsense module	7
18650 Lion Battery	7
LSM9DS1 IMU	7
10k Resistor	3
Battery holder	5
Push button	3
Force Sensitive Resistor	3
Solder	3

lead time or are less reliable. In addition, having additional spares helps reduce the immediate impact of parts either going out of stock for long periods or parts no longer being produced.

In the RAM analysis we considered mean time to repair assuming the components required were already available and someone was available to fix them immediately. Neither of these cases are true in our context. It can take 3-7 business days for some of the parts within the T-Stick to arrive in the lab. In addition it is rare that someone is immediately available to fix a T-Stick once it is broken. Table 4 shows the lead time is shown for each component.

To optimise for the minimum amount of spare parts required to maintain our availability target we will use a Monte Carlo simulation of the system. We will start with 5 of each component and no spare T-Sticks. We will simulate a period of 50000 performance sessions lasting 8 hours each. During each time step of the simulation we will check if a part has failed and if a technician is available to fix it. To simulate technician availability we consider a worker who is available 10 hours to fix T-Sticks. During each hour they fix a work on a single part. Once 10 hours have past they are not available again for another 5 performance sessions (40 hours).

The operational availability is computed by dividing taking the minimum working time from the 3 T-Sticks and dividing that by the total simulation time. this average is computed every 8 hours and saved. If the average availability is not within a 95% confidence interval of the achieved availability we will increase the amount of the component that we had to order the most by 1.

5.3 Results and Discussion

We analysed the maximum availability that can be reached with no spare T-Sticks. Given no spare T-Sticks the availability was around $80.83 \pm 2.31\%$. This is far below the 99% target we are aiming for and the 96.94% inherent availability calculated earlier. Increasing the spare amounts of other components did not significantly improve the availability beyond 82%.

As the number of spare T-Sticks is increased we see that the operational availability slowly approached 99%. Typically it required 3 spare T-Sticks for the average availability to reach 99%. These results are shown in table 5.

These results are a bit troubling as it indicates that we need more spares than T-Sticks in use to maintain a high availability target. This can cause issues when running longitudinal studies with many participants or performances with a large amount of performers. Running this analysis with 10 T-Sticks in use shows that we need approximately 6 additional T-Sticks to maintain our availability target of 99%.

These results are due to the lead time introduced by the

Table 5: Simulated Results for the Operational Availability of the T-Stick

Iteration	Spare T-Sticks	Availability (%)
1	0	$80.31 \pm 2.31\%$
2	1	$92.11 \pm 0.81\%$
3	2	$97.21 \pm 0.83\%$
4	3	$99.32 \pm 0.34\%$

lack of reliable technician availability. If a technician could be available for every performance session the number of spare T-Sticks needed drops to 1. However, it may not be a realistic option for a research organisation to keep a technician hired year round just for maintenance of their instruments. Therefore, significantly higher amounts of spare instruments are needed to maintain high levels of availability.

To maintain our target of 99% availability for 3 T-Sticks it requires spending 1659.75 CAD initially for all the spare components and spare T-Sticks. Over the entire run time we spent 16,500 CAD to maintain the instrument which is about 0.33 CAD per session. This cost only includes the cost of the components, it does not include the labor cost of the technician. This number did not fluctuate heavily when changing the amount of spare T-Sticks. The cost per session stayed around 0.29 - 0.33 CAD.

6. CONCLUSION AND FUTURE WORK

In this paper we discussed maintenance and support of digital musical instruments using the T-Stick as a case study. A brief RAM analysis is conducted on the T-Stick to estimates its reliability characteristics, in particular the mean time between failure, mean time to repair and availability of the T-Stick. Spare parts and instruments are required to keep a high level of availability of the instrument for performances and experiments. For the current batch of Sopranino T-Sticks, it was found that to maintain a 99% availability target for 3 T-Sticks we need to have 3 spare T-Sticks.

Future works include extending this analysis to consider signal accuracy and precision requirements, as well as the impacts of the operating conditions on the T-Stick. This includes the impact of the venue, amounts of devices on the same network as the T-Stick, and the quality of the Wi-Fi signal. More in-depth testing on the T-Stick will be conducted to more accurately estimate the failure rates of the T-Stick's subsystems. We plan to use this analysis to contribute towards a new iteration of the T-Stick that improves on the reliability and maintainability of the instrument.

As shown with the analysis on the T-Stick this analysis can help with understanding the performance of your instrument under different conditions and the potential maintenance cost of the instrument. Instrument designers looking to transition their instruments away from research prototypes to stable instruments can use this technique to understand the current limitations of the instrument and work towards improving the reliability of their instruments.

Maintenance and support is an ongoing process. The techniques presented in this article provide different views on how to measure and understand reliability and maintenance. As we learned with our experience building T-Sticks and conducting this analysis, poor reliability can come from multiple sources such as poor design, low inherent reliability of components and poor manufacturing. Understanding where the reliability issues from your instrument is coming from and how it is impacting the instrument is key for improving the reliability of your instrument.

7. ACKNOWLEDGMENTS

This work was possible thanks to the IDMIL, and CIRMMT. This research was also partially supported by a Discovery grant from the Natural Sciences and Engineering Council of Canada.

8. ETHICS STATEMENT

There are no observed conflicts of interest. All researchers participated consensually in the activities described in this document.

9. REFERENCES

[1] T. I. Băjenescu and M. I. Bâzu. *Introduction*, pages 1–42. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.

[2] S. Chu, F. Liu, and Z. Wei. The Study on Dispatch Reliability Prediction Model of Civil Aircraft. *The Open Mechanical Engineering Journal*, 8:828–832, 12 2014.

[3] P. R. Cook. Principles for designing computer music controllers. In *Proceedings of the CHI'01 Workshop on New Interfaces for Musical Expression (NIME-01)*, pages 3–6, 2001.

[4] Department of Defense. Military handbook: Reliability prediction of electronic equipment: Mil-hdbk-217f, December 1991.

[5] V. Goudard. Ephemeral instruments. In *Proc. International Conference on New Interfaces for Musical Expression (NIME)*, pages 349–354, 2019.

[6] INCOSE. *Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities.* John Wiley and Sons, 4 edition, 2015.

[7] J. Lienig and H. Bruemmer. *Reliability Analysis*, pages 45–73. Springer International Publishing, Cham, 2017.

[8] J. Malloch, S. Sinclair, and M. Wanderley. Distributed Tools for Interactive Design of Heterogeneous Signal Networks. pages 1–25, 2014.

[9] J. Malloch and M. M. Wanderley. The T-Stick: From Musical Interface to Musical Instrument. In *Proc. International Conference on New Interfaces for Musical Expression (NIME)*, pages 66–69, 2007.

[10] C. B. Medeiros and M. M. Wanderley. A Comprehensive Review of Sensors and Instrumentation Methods in Devices for Musical Expression. *Sensors*, 14(8):13556–13591, 2014.

[11] F. Morreale and A. P. McPherson. Design for Longevity: Ongoing Use of Instruments from NIME 2010-14. In *Proc. International Conference on New Interfaces for Musical Expression (NIME)*, pages 192–197, 2017.

[12] F. Morreale, A. P. McPherson, and M. M. Wanderley. NIME Identity from the Performer's Perspective. In *Proc. International Conference on New Interfaces for Musical Expression (NIME)*, pages 168–173, 2018.

[13] SEBoK. System Reliability, Availability, and Maintainability — SEBoK, 2021. [Online; accessed 1-February-2023].

[14] Siemens AG Munich. SN 29500. *Failure Rates for Electronics and Electromechanical Components*, 2022.

[15] M. M. Wanderley and N. Orio. Evaluation of input devices for musical expression: Borrowing tools from HCI. 26(3):62–76, 2002.