

Master thesis : Head phantoms: Design and implementation of a physiological signal generator on a PCB for the validation of new polysomnography systems

Auteur : Boniver, Julien

Promoteur(s) : Redouté, Jean-Michel

Faculté : Faculté des Sciences appliquées

Diplôme : Master : ingénieur civil électricien, à finalité spécialisée en Neuromorphic Engineering

Année académique : 2025-2026

URI/URL : <http://hdl.handle.net/2268.2/26100>

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.

Master's Thesis presented to obtain the degree of:
**Master of Science for "Neuromorphic Engineering" in
Electrical Engineering**



Head phantoms: Design and implementation of a physiological signal generator on a PCB for the validation of new polysomnography systems

Performed by:
Julien BONIVER (s212020)

Supervisor:
Prof. Jean-Michel REDOUTÉ

Daily supervisor:
Laura TOUSSAINT

Master's Thesis Committee

- Prof. Jean-Michel Redouté (Université de Liège), *Advisor*
- Prof. Philippe Vanderbemben (Université de Liège)
- Prof. Pierre Sacré (Université de Liège)

Declaration on the Use of Automatic Tools for Writing the Manuscript

I hereby certify that I have not used any generative intelligence tool in the writing of text, graphics, images, or data reproduced in this manuscript.

The writing tools used in this manuscript are the following :

- **DeepL Translator:** All along the project.
- **Reverso:** All along the project.
- **Grammarly:** All along the project.

Abstract

Polysomnography (PSG) is the gold standard for diagnosing sleep disorders and relies on the simultaneous overnight acquisition of several physiological signals. Validating new PSG acquisition systems before they reach patients requires a testbench capable of reproducing realistic multimodal signals under controlled conditions; yet existing solutions either target a single modality, lack configurable artifact injection, or remain inaccessible because of cost.

This work presents the design, implementation, and characterization of an accessible, multimodal physiological signal generator on a printed circuit board for the validation of new PSG systems. The signals of interest in the scope of this work are EEG, EOG, ECG, and EMG. The device replays real recordings from the Haaglanden Medisch Centrum sleep staging database through eight channels: seven single-ended (3 EEG, 2 EOG, EMG) and one differential (ECG). Cardiac and ocular artifacts are first extracted through a method-per-artifact pipeline, allowing them to be re-injected at user-defined gains. The hardware combines a 12-bit eight-channel DAC, a digitally controlled two-stage resistive divider sized per modality, and precision buffers. Three playback modes (raw replay, artifact mixing, and on-demand artifact injection), each with optional sleep-stage selection, are controllable from a PC terminal without reflashing. Signals are stored on an SD card, and dry-electrode connection is supported through ten on-board conductive pads. The complete device costs approximately 200 euros.

The generator presents an output range from 62 μV to 79 mV with a best-case noise floor near 10 μV , a faithful reproduction on the ECG differential channel across all configurations. On the single-ended channels, inter-channel coupling through the shared analog reference node limits faithful simultaneous multichannel replay below the millivolt range; the mechanism is analyzed and mitigation strategies are proposed. The supply noise is shown to be rejected by the differential output topology, making the generator effectively independent of the choice of power source. Overall, the work demonstrates that accessible, multimodal, hardware-level PSG validation with configurable artifact injection is achievable and provides a documented basis for future versions.

Acknowledgements

I would first like to express my sincere gratitude to my supervisor, Prof. Jean-Michel Redouté, for giving me the opportunity to work on a subject that suited my interests, and for his support and trust throughout this thesis.

I would like to sincerely thank Laura Toussaint, PhD student, for her ongoing support and valuable advice at every step of my thesis. Her availability and insightful feedback were greatly appreciated throughout the development of this work.

I would like to warmly thank the entire Microsys team for their hospitality and support during this project; in particular, Gabriel Digregorio for providing me with technical resources and valuable advice, Angel Calderon Jimenez for his help with the hardware aspects, and Morgan Diepart for his expertise in hardware and firmware debbuging.

Finally, I would like to express my heartfelt gratitude to my family for their support and trust throughout my studies, and to everyone who has been part of this chapter.

Contents

Introduction	1
1 State of the art	2
1 Polysomnography and physiological signals	2
2 Artifacts in PSG signals	6
2.1 External artifacts	6
2.2 Internal artifacts	7
3 Review of existing simulators	9
3.1 Head phantoms	9
3.2 Physiological signal generators	11
3.2.1 Commercial patient simulators	12
3.2.2 Academic hardware generators	13
3.2.3 Software-only biosignal generators	13
3.3 Artifacts injection in physiological signal generators	14
3.4 Synthesis	15
2 Design of the physiological signal generator	16
1 Signal processing	16
1.1 Data preprocessing	17
1.2 Internal artifacts removal	17
1.2.1 ECG artifacts removal	17
1.2.2 EOG artifacts removal	18
1.2.3 EMG artifacts removal	19
1.2.4 Validation metrics	20
1.3 Data export	21
1.4 Artifacts injection	21
1.4.1 Artifacts mixing	22
1.4.2 Artifacts library	22
2 Hardware	23
2.1 Circuit design	23
2.1.1 Power supply and analog reference	24
2.1.2 Digital part	25
2.1.3 Analog part	27
2.1.4 Noise	32
2.1.5 Simulation results	36
2.2 Printed circuit board design	41
2.2.1 Layer stack-up	41
2.2.2 Components placement	42

2.2.3	Interconnects to the Margherita board	42
2.2.4	SPI daisy chain routing	43
2.2.5	V_{ref} node	44
2.2.6	ECG differential output	45
2.2.7	Identified design omissions and post fabrication fixes	46
3	Firmware	47
3.1	Firmware architecture	47
3.2	Schedulability	48
3.3	Playback modes	49
3.3.1	Raw signal playback mode	49
3.3.2	Artifacts mixing playback mode	50
3.3.3	Artifacts injection mode	50
3.3.4	Specific sleep stage playback option	51
3	Results and discussion	52
1	Measurement method and protocol	52
1.1	Characterization protocol	52
1.1.1	DC characterization	52
1.1.2	Dynamic characterization	55
2	DC characterization results	58
2.1	Output range	58
2.2	Output linearity	60
2.3	Noise floor of the generator	64
2.4	Discussion about the DC characteristics of the generator	65
3	Dynamic characterization results	67
3.1	Supply noise	67
3.2	Physiological signal playback	68
3.3	Validation of the Generator Modes	71
3.3.1	Artifacts mixing	72
3.3.2	Artifacts injection	72
3.3.3	Sleep stage playback	73
3.4	Discussion about the dynamic characteristics of the generator	75
4	Conclusion and perspectives	77
1	Conclusion	77
2	Perspectives	78
	Appendices	79
A	Generator schematics	79
B	Generator PCB layout	85
C	Margerhita board	86
D	User interface	87
	Bibliography	88

Introduction

Sleep is a vital need for every human being, and yet not every person is equal in terms of sleep quality. About 9% to 24% of the population suffer from sleep apnea, while 2% to 4% are affected by obstructive sleep apnea (OSA) syndrome [1]. This disorder is commonly underdiagnosed and impacts the patient's quality of life and health, and is considered an important public health problem [2]. OSA is diagnosed using polysomnography (PSG), which consists of recording several physiological signals from the patient over a full night of sleep at the hospital, using a setup such as the one shown in Figure 1 [3]. PSG is the gold standard for diagnosing sleep disorders, including OSA and many others. The acquisition systems used for these recordings are constantly evolving, as patient comfort and recording quality can always be improved. To develop and validate such new devices before using them on real patients, a testbench is needed to replay real physiological signals in a configurable environment. Two types of simulators exist for testing such monitoring systems: head phantoms and physiological signal generators. However, none of them simultaneously provides the replay of multimodal physiological signal recordings, configurable artifact injection, and accessibility.




Figure 1: Polysomnography setup for the recordings of physiological signals.

The goal of this work is to design and implement a physiological signal generator for testing and validating new PSG monitoring devices. The device is intended to be an accessible, multimodal hardware generator that replays real PSG signals, injects configurable artifacts, supports the connection of dry electrodes, and can be controlled from a PC terminal without reflashing.

The remainder of this work is organized as follows. Chapter 1 introduces PSG and the associated signals, and reviews the existing simulators, which together motivate this work. Chapter 2 describes the design of the generator and is divided into three sections. The first describes the signal processing applied to recorded PSG signals to separate them from their artifacts, a prerequisite for controllable artifact injection. The second covers the hardware design of the generator board, which is the core of this work. The third describes the firmware of the generator. Chapter 3 presents the methods used to characterize the generator and reports and discusses the results. The work closes with a conclusion that addresses the limitations of the generator and perspectives for future versions.

Chapter 1

State of the art

When developing biomedical devices that will be in contact with human patients, the testing and validation process cannot be done directly on human subjects. It is therefore necessary to develop devices that will be used to test and validate biomedical devices. Such test and validation devices that simulate human subjects are usually called phantoms. These phantoms have to be as representative as possible to a real human subject in order to validate biomedical devices. The goal of this work is to design a bio-signal generator in order to test and validate new biomedical devices. More particularly, the generator is focused on the validation of new polysomnography monitoring devices. Therefore, it is important to have some background knowledge about polysomnography and the signals of interest in this biomedical diagnosis. The goal of this section is to describe the physiological signals monitored during polysomnography, how they are acquired, what their electrical properties are, what influences them, and to review the phantoms and other devices currently used to validate PSG monitoring systems.

1 Polysomnography and physiological signals

Polysomnography (PSG): PSG is an overnight process used to collect several physiological parameters during sleep. PSG is considered to be the gold standard for diagnosing sleep-related disorders such as obstructive sleep apnea, central sleep apnea, hypoxia, narcolepsy and rapid eye movement sleep behavior disorder [4]. PSG involves the acquisition of precise physiological signals: electroencephalogram (EEG), an electrooculogram (EOG), a chin and leg electromyogram (EMG), an electrocardiogram (ECG), and pulse oximetry, as well as airflow and respiratory effort [4]. Thanks to standard criteria, this combination of physiological signals allows for the characterization of the sleep stages : Wake, N1, N2, N3, and REM [5]. The generator developed in this work supports selective sleep-stage playback, which can be used to validate sleep-stage classification algorithms, although the sleep stages themselves are not described in detail in this work. The AASM [5] recommends technical and digital specifications for the different signals in order to be recognized as true PSG recordings. For all signals, the minimum recording sampling rate should be 200 Hz and the recommended sampling rate is 500 Hz. The recommendation regarding sampling frequency is not arbitrary: the Nyquist-Shannon theorem states that the sampling frequency of a given signal must be at least twice the signal's frequency in order to ensure an accurate reconstruction of the signal. The useful bandwidth of the signals ranges from a few hertz to 100 Hz; Nyquist's theorem is therefore satisfied. The remaining specifications, along with a

description of the signals, are presented later in this section.

The electroencephalogram (EEG): The EEG records principally the electrical activity of the cortex and also reflects deeper brain activity. The EEG signal is a differential potential between two electrodes [6].

There are many possible electrode locations on the scalp for recording EEG signals, as described in Figure 1.1 [6]. For PSG recordings, the recommended EEG derivations are F4-M1, C4-M1, and O2-M1 (Fz-Cz, Cz-Oz, and C4-M1 are also acceptable) in the 10-20 electrode placement system [5], as represented in Figure 1.1. Different wave patterns in the EEG signal during wakefulness and sleep are useful for differentiating sleep stages. The activity is separated into four frequency bands: alpha, beta, theta, and delta waves. Beta waves are defined between 13 and 25 Hz, mostly originating in the frontal and central regions, and are present during wakefulness and drowsiness. Alpha waves, known as the posterior dominant rhythm, are defined between 8 and 13 Hz, are best seen in the occipital leads, and are present during relaxed wakefulness and central nervous system arousal. Theta waves are defined between 4 and 8 Hz, originate from the central vertex region, and are the most common waves in the EEG recording. Finally, delta waves are defined between 0.5 and 2 Hz, are most often seen in the frontal region, and characterize stage N3 of sleep. Besides these four frequency bands, two other wave patterns help to characterize sleep stages: sleep spindles (11–16 Hz, in the central vertex) and K-complexes (sharp slow waves in the central vertex), both characteristic of stage N2 sleep [6]. The useful EEG frequency band ranges from 0.5 to 30 Hz; the acquisition system is therefore recommended to apply a band-pass filter with a low cut-off at 0.3 Hz and a high cut-off at 35 Hz. This frequency band is the useful band for the generated EEG signals. The EEG signal amplitudes range from a few tens of μV to more than 75 μV depending on the sleep stage. The generator needs an ultra-low output amplitude, down to a few tens of μV , in order to faithfully reproduce the amplitude range of the EEG.




Figure 1.1: AASM recommended electrode locations of International 10-20 system for EEG recording.

The electrooculogram (EOG): The EOG records eye movements. It is based on the corneo-retinal electrical dipole, whose potential variations are detected by surface electrodes [7].

Two derivations are recommended: E1-M2, placed 1 cm below the outer canthus of the left eye, and E2-M2, placed 1 cm above the outer canthus of the right eye (M2 being the right mastoid electrode location shown in Figure 1.1) [5], as shown in Figure 1.2[6]. Any eye movement causes the cornea to shift toward one electrode while the retina shifts away from it, generating a measurable potential difference. In the absence of movement, the relative position between the eye and the electrodes remains unchanged, so the EOG leads record no signal. The EOG plays a key role in distinguishing the rapid eye movement (REM) stage from the other sleep stages. REM consists of sharp conjugate saccades with a duration of less than 0.5 seconds, whereas slow rolling eye movements (SREMs) are conjugate, often rhythmic, and last more than 0.5 seconds [6]. The EOG is a very low-frequency signal, SREMs have frequencies below 2 Hz and REMs range between 2 and 10 Hz, so the useful frequency band extends from near DC to 10 Hz. In practice, the acquisition filter settings should include a high-pass filter with a cut-off at 0.3 Hz and a low-pass filter with a cut-off at 35 Hz, which defines the useful frequency band for the generated EOG signals. EOG amplitudes are greater than those of the EEG and range from $50\text{ }\mu\text{V}$ to several hundred μV , large blinks can even reach several mV. The EOG channel must therefore be capable of handling a range from approximately $100\text{ }\mu\text{V}$ to a few mV.




Figure 1.2: AASM recommended electrode locations for EOG recording.

The chin electromyogram (EMG): The EMG measures muscle tone. Usually, three electrodes are used to record the chin EMG (Figure 1.3 adapted from [8]): one mental electrode serves as the reference, and two submental electrodes detect the activity of the floor of the mouth as a marker of muscle atonia in REM sleep [4]. In practice, only one of the two submental electrodes actively records data, while the other is used as a backup [5]. EMG activity decreases during sleep and is minimal during REM sleep [6]. The frequency band of the EMG is wider than those of the other signals, ranging from 10 to 500 Hz; however, the relevant signal content lies between 10 and 100 Hz. The acquisition filter settings include a high-pass filter with a cut-off at 10 Hz and a low-pass filter with a cut-off at 100 Hz, which defines the useful frequency band for the generated EMG signals. The amplitude of the EMG signal also covers a wide range, from a few tens of μV to a few mV peak-to-peak, depending on muscle contraction. The EMG channel must therefore be capable of handling a range from a few tens of μV to a few mV.




Figure 1.3: Recommended electrode locations for chin EMG recording

The electrocardiogram (ECG): The ECG records the heart’s electrical activity. Several methods exist for recording the ECG, using 3, 5, or up to 12 electrodes. The simplest configuration consists of three electrodes: a bipolar signal acquisition using two active electrodes and one reference electrode. Several derivations are possible with three electrodes, depending on where the electrodes are located and what they are used for. These include the lead I, II, and III ECG derivations [9]. For PSG recordings, the standard derivation is a modified lead II [4], obtained by placing one active electrode below the right clavicle and the other active electrode over the left lower thorax, as shown in Figure 1.4a [10].

The ECG signal is characterized by five waves corresponding to the heart’s activity. The P-wave is the first positive deflection on the ECG and corresponds to atrial depolarization. The QRS complex corresponds to atrial repolarization and to the depolarization of the right and left ventricles. The T-wave arises from the repolarization of the ventricles [11]. An example of the ECG waveform is shown in Figure 1.4b [11]. The QRS complex occupies a frequency band from 5 to 25 Hz and represents the most useful information. The acquisition filter settings include a high-pass filter with a cut-off at 0.3 Hz and a low-pass filter with a cut-off at 70 Hz, which defines the useful frequency band for the generated ECG signals. ECG amplitudes range from several hundred μV up to a few mV peak-to-peak. The ECG has the highest amplitude in the PSG setup, which explains why it interferes with the other channels. This is one of the reasons the artifact injection pipeline treats it as a separate stream.




Figure 1.4: a) Electrode locations for lead II configuration for ECG recording. b) Classical ECG waveform.

Other PSG signals: While other signals are recorded during PSG, they are described here only for context and fall outside the scope of this work.

Oxygen saturation (SpO_2) is recorded using pulse oximetry to help with hypopnea scoring. The leg electromyogram (right and left anterior tibialis EMG) involves placing electrodes on the left and right anterior tibialis muscles to detect periodic limb movement disorder (PLMD). Nasal and oral airflow are recorded using an oronasal thermal sensor that detects temperature changes during expiration and inspiration; this signal is used to detect apneas. Respiratory effort is recorded at the thorax and abdomen using two respiratory inductance plethysmography (RIP) sensors. These signals are used to detect different respiratory events, including obstructive apneas, obstructive hypopneas, and central apneas. The patient’s position can also be video-recorded during the night in order to correlate apneas with certain sleep positions

In the framework of this master’s thesis, the EEG, EOG, chin EMG, and ECG are the signals of interest, as they are the electrical signals that will be produced by the physiological signal generator. One goal of the generator is to replay these signals at the physiological amplitudes described above with high fidelity to the original signal.

2 Artifacts in PSG signals

The signals recorded during a PSG test are permanently contaminated by electrical activity external to the channel of interest. These signals operate at a scale of about ten microvolts to a few millivolts. At this level, contamination is rarely negligible; it can equal or exceed the useful signal in amplitude. Consequently, it risks masking the events that the recording is meant to capture, exceeding the instrument range, and generating incorrect output or false alarms in downstream automated analysis[12]. The artifact-removal literature classifies these contaminations into two categories, external and internal artifacts, which are distinguished by their origin [13], [14]. External artifacts are introduced by the measurement chain itself or by its electromagnetic environment. Internal artifacts are generated by the patient’s own body, but originate from a different physiological source than the one targeted by the channel. The distinction matters because the two categories call for fundamentally different mitigation strategies. Internal artifacts are correlated with a physiological reference that is usually available on another channel of the same recording and can therefore be modeled and subtracted. External artifacts have no such reference on the patient and are addressed primarily by analog front-end design, shielding, electrode preparation, and filtering. The remainder of this section reviews each family in turn.

2.1 External artifacts

External artifacts are introduced by the measurement chain or by the environment, and the artifact-removal literature classifies them as extrinsic, distinct from the intrinsic contamination produced by the subject’s own body [14]. They are uncorrelated with the subject’s physiology, and no reference channel taken on the patient can predict them, which rules out the subtraction-based methods that work for internal artifacts and shifts the mitigation strategies toward analog design, electrode preparation, and filtering [14]. They can be separated in two main categories.

Environmental artifacts: Any electrical system in a normal environment is subject to parasitic signals. These are induced by AC fields and can be classified into two categories: **electric field** coupling, which induces currents in conductive materials, and **magnetic field** coupling, which induces an electromotive force in conductive loops. PSG consists of

measuring small-amplitude signals generated by the human body using electrodes wired to an acquisition system. Therefore, both phenomena directly influence the measurement [15]. Magnetic field coupling can be mitigated by making the loop area formed by the connection cables as small as possible. Electric field coupling has two effects: displacement currents into the leads and into the human body. The first induces interferences in the signal if the currents are unequal between the leads, which can result from electrode impedance mismatch; this effect can be mitigated with proper electrode shielding. The second induces interferences in the signal as a parasitic current flows through the body when it is grounded. The result is that parts of the body are at different potentials, and the potential difference is amplified by the acquisition system. Adequate placement of the ground electrode can mitigate this effect [15]. In PSG, the 50 Hz AC field falls inside the EMG and ECG bands, just above the upper EEG band, and well outside the EOG band. Even when precautions are taken in the setup, some interferences remain, which is why they are typically attenuated by a narrow notch filter applied selectively to the channels where they fall in-band [16], or by a low-pass filter for the channels where they fall out-of-band, as recommended in the PSG settings recommendations described above.

Motion and electrode/contact artifacts: Every mechanical disturbance of the electrode–skin system, such as body movement, head repositioning on the pillow, cable sway, or progressive changes in the interface due to sweating or gel drying, modulates the electrical contact and introduces distortion to the recorded signals [17]. Tam and Webster [12] showed that, when surface electrodes and conducting paste are used, a major cause of this contamination is not a mechanical displacement of the electrode itself but a change in the skin potential at the electrode–skin interface. The contamination does not have a specific frequency band and often overlaps the frequency band of the signal, which eliminates the possibility of using classical filtering approaches [17]. Symeonidou et al. [18] further demonstrated, using a gelatin head phantom moved by a robotic platform, that cable sway alone is a major contributor at low motion frequencies.

None of these categories is explicitly reproduced by the physiological signal generator developed in this work. The generator is a wired electronic system operating in the same environment as a patient in a PSG and is therefore subject to the same phenomena. The generator used in an AC field environment experiences electric field coupling, and the measurement chain experiences magnetic coupling, cable sway, and electrode-pad disturbance. Any external artifact present at the input of the device under test is consequently reproduced naturally by the measurement chain itself rather than synthesized at the source.

2.2 Internal artifacts

A common framework in clinical neurophysiology distinguishes **near-field** from **far-field** potentials [19]: A near-field potential is recorded near the active source, and its amplitude decreases sharply with the distance between the source and the electrode, whereas a far-field potential is recorded at a greater distance and reaches the electrode via conduction through the body, with a much weaker dependence on the electrode’s position. The transition between the two regimes is governed by the geometry of the electrode–source system, formalized by the concept of the lead field: the spatial weighting by which a given electrode integrates the contributions of each dipole source distributed throughout the body [20]. Each electrode is particularly sensitive to sources located within the high-intensity region of its near-field, that is, those that generate a near-field potential, and becomes progressively less sensitive to more distant sources, which contribute in the form of far-field potentials. Internal artifacts occur

when an electrode’s derivation field overlaps with a bioelectric source other than the one the channel is intended to record. Because the human body is a single conductive volume, every active source, the heart, the eyes, the muscles, contributes simultaneously to every electrode of the montage. What changes from one channel to another is the distance to each source, and therefore the regime in which the contamination is picked up. The same source can thus produce a near-field contamination on one channel and a far-field contamination on another. The ocular dipole, for example, is recorded as the intended near-field signal of an EOG channel, contaminates frontal EEG electrodes in a still near-field regime where amplitude depends strongly on electrode position, and reaches occipital EEG electrodes as a far-field contribution. Three sources dominate in PSG: cardiac, ocular and muscular activity.

Cardiac field artifact: The heart is a powerful bioelectric generator in the body; its electric field, on the order of 1 to 2 mV in amplitude [21], propagates throughout the body via volume conduction and is therefore picked up by all electrodes in a PSG setup [22]. On the EEG, EOG and chin EMG channels, the heart lies far outside the high-magnitude region of the lead field, and the cardiac contamination appears as a true far-field potential, small in amplitude, smoothly distributed across electrodes, and time-locked to each cardiac cycle. It is most prominent during the QRS complex and the T wave, and shows a topographic asymmetry between the two hemispheres [22].

Ocular artifacts: The eyeball behaves as a permanent electrical dipole, with the cornea positively charged with respect to the retina (the corneo-retinal dipole) [23]. Any rotation of the eye, saccades, slow rolling movements, smooth pursuit, changes the orientation of this dipole in the head and produces a potential change at every electrode of the head montage. Because the dipole is small and stationary in position, the amplitude of the resulting potential declines steeply, approximately inversely proportional to the square of the distance between the electrode and the eye [23]. This distance dependence means that ocular contamination varies across the scalp. It is largest on the EOG electrodes themselves, where the dipole is, in fact, the intended near-field source of the channel. It remains in a near-field regime on the frontal EEG electrodes, where amplitudes depend strongly on electrode position and can exceed the background brain activity. On central, parietal, and occipital electrodes, the contamination progressively shifts toward a far-field signature, with weaker position dependence [23]. Blinks add a second mechanism, as the eyelid slides over the positively charged cornea, it generates a positive potential on the EEG that is independent of gaze direction [24]. Most of the spectral energy of ocular artifacts lies below 4 Hz [25], which falls directly within the delta band relevant to sleep scoring.

Muscle artifacts: Each contracting motor unit acts as a small bioelectric source whose field propagates through the surrounding tissue [26]. The muscles that contaminate the EEG and EOG channels are the frontalis, temporalis, and occipitalis muscles [27]. The chin EMG electrodes are immediately adjacent to the underlying muscle, so they sit inside the high-magnitude region of the lead field of every active motor unit beneath them; therefore, the muscle contamination is picked up in the near-field regime. Three consequences follow. First, the contamination can reach large amplitudes, often higher than the signal of interest [27]. Second, it falls off steeply with distance from the active muscle, so the contamination is highly localized in space rather than being smoothly distributed across the montage. Third, it is no longer time-locked to a single physiological event but reflects the irregular firing pattern of the underlying muscles. In addition, surface EMG occupies a broad spectral band

that overlaps the upper EEG bands (beta and gamma) [28], which prevents separation by simple bandpass filtering without sacrificing legitimate brain signals.

It is important to understand how the PSG montage interacts with these properties. The muscles that contaminate the EEG derivations are the frontalis, temporalis, and occipitalis muscles, whereas the only EMG channel recorded in the standard montage is the submental derivation. These are anatomically distinct motor groups, recruited by different reflex and voluntary drives. The chin EMG and the scalp EMG contamination are only linked at certain times; both rise during arousals and movements and fall during consolidated sleep. Therefore, the chin EMG channel cannot serve as a faithful reference for the muscle activity that contaminates the EEG. The multi-channel separation techniques developed for scalp EMG removal in the Brain Computer Interface (BCI) and high-density EEG literature [25] are not well-conditioned for a four-channel EEG montage. The consequences of this for the artifact-removal pipeline and the simulator architecture are discussed in Section 1.2.3.

The internal artifact family is the one of interest in this work. The methods used to remove each of these three artifact families from the data and store them as a separate stream are described in Section 1.2.

3 Review of existing simulators

The development of new PSG monitoring devices, aimed at improving patient comfort and diagnostic accuracy, has driven a parallel need for reliable validation tools. Before any new sleep-monitoring device can be tested on human subjects, its behavior must first be characterized under reproducible conditions. Two main categories of simulators are commonly used for this purpose: head phantoms and biosignal generators. Understanding what each simulator represents and how they work is an essential prerequisite to the design of a physiological signal generator.

3.1 Head phantoms

Head phantoms are physical models that reproduce the electrical and geometric properties of the human head, providing a controlled medium for the validation of EEG electrodes and acquisition systems [29]. Their role is to recreate the path that a biopotential follows from its source to the recording electrode. The head is an inhomogeneous volume conductor, composed principally of the brain, skull, and scalp, which have different conductivities. The scalp and grey matter of the brain have approximately the same conductivity, around 0.41 S/m, while the skull has around 0.02 S/m [30]. Reproducing the different conductivity regions with stability over time and accessibility remains a challenge for phantom design.

In the course of developing a low-cost 3D-printed phantom, Akor et al. [31] identify seven categories of phantom technologies: commercial injection-molded, saline and electrolyte solutions, gelatin and hydrogel, 3D-printed conductive, silicone and elastomer, multi-material, and emerging textile-based phantoms. Examples of several of these categories are shown in Figure 1.5 [31, 32, 33, 34].

Commercial injection-molded phantoms (1.5a) are the gold standard for regulatory testing, providing exceptional consistency and multi-year durability with a resistance through the head of 10 Ω cm to 20 Ω cm. Built around standardized anatomical models, they incorporate internal electrodes capable of generating calibrated signals such as EEG. However, their

tooling costs several thousand euros, and the unit price is several hundred euros, limiting access to this equipment to laboratories with substantial financial resources [31].

Lower-cost alternatives trade durability for accessibility. Saline phantoms offer an easily tunable, tissue-like conductivity set by the electrolyte concentration, but their useful life is only a few hours. Evaporation progressively raises the concentration and shifts the conductivity, contamination compromises repeatability, and the solution accelerates electrode oxidation, making them unsuitable for extended testing sessions. Gelatin phantoms (Figure 1.5b) last longer, on the order of a few weeks [31], but degrade through dehydration and microbial growth, exhibit drifting mechanical properties, and are difficult to reproduce with consistent electrical characteristics from batch to batch. Their electrical behavior is nonetheless well characterized. Owda and Casson [33] measured several gelatin formulations and showed that the phantom reproduces the body-electrode frequency response over the 20-1000 Hz band. They also showed that the addition of NaCl tunes the impedance profile to match a chosen target tissue in a repeatable manner, and that the contact-impedance and group-delay response remain stable over seven days. Despite their limitations, saline and gelatin phantoms remain well suited to educational use, electrical-property evaluation, and exploratory research.

More recent approaches, including 3D-printed conductive phantoms (Figure 1.5c), silicone-based models, and textile-based hybrids (Figure 1.5d), partially address these limitations while preserving the electrical realism of the categories above. 3D-printed phantoms using conductive filament achieve a tunable skull resistivity, typically between $15\,\Omega\text{cm}$ and $100\,\Omega\text{cm}$, by varying the infill percentage of the filament, which provides durability and stability of the phantom’s electrical properties. In a parallel direction, Tseghai et al. [34] developed long-lasting textile phantoms that offer durability and stability of electrical characteristics, accurately reproducing the human-body electrode frequency response and yielding a signal-to-noise ratio of $15.95 \pm 1.67\text{ dB}$, higher than that of a ballistic-gelatin phantom [34]. Silicone and elastomer models pursue durability through a third material route, typically providing a resistivity between $5\,\Omega \cdot \text{cm}$ and $15\,\Omega \cdot \text{cm}$.

Beyond electrode and impedance testing, anatomically accurate phantoms have also been used to validate effective connectivity measures by combining physical models with neural mass simulations [35], and to systematically characterize motion-induced artifacts in mobile EEG recordings [36]. Table 1.1 summarizes the electrical characteristics, durability, and cost of the principal phantom categories, adapted from [31].




Figure 1.5: a) Injection molded head phantom. b) Gelatin head phantom. c) 3-D head phantom. d) textile head phantom

Table 1.1: Comparison of head phantom technologies for EEG/PSG validation.

Type	Cost (€)	Fabrication time	Resistivity	Key trade-off
Commercial	350–600	Weeks	10–20 Ωcm	Consistency and durability but costly
Saline	10–40	Hours	Tunable	Accessibility but electrodes corrosion and short shelf-life
Hydrogel	30–90	1–2 d	2–5 Ωcm	Accessibility but short shelf-life
3D-printed	40–180	1–3 d	15–100 Ωcm	Flexible design but lacks validation
Silicone	100–350	2–3 d	5–15 Ωcm	mechanical fidelity but complex fabrication
Multi-material	200–700+	3–7 d	Heterogeneous	Most realistic but very limited access
Textile	90–200	1–2 d	Not reported	Lightweight and durability but complexity in electrical property control

3.2 Physiological signal generators

While head phantoms are well established for electrode-level validation, assessing contact impedance, spatial signal pickup, and the electrode-tissue interface, their utility for comprehensive PSG device validation remains limited. Phantoms address the question of where a biosignal is acquired, providing a controlled conductive medium between source and electrode, but they do not address the question of what signal is being acquired. They offer

little control over the temporal and spectral content of the signals reaching the acquisition system, and they do not, by themselves, reproduce the multiple, time-aligned physiological signals (EEG, EOG, EMG, and ECG) that PSG devices must reliably capture across distinct frequency bands and amplitude ranges, nor the artifacts that contaminate clinical recordings. Head phantoms are therefore best understood as complementary to physiological signal generators rather than substitutes. Phantoms model where signals propagate, whereas physiological signal generators model what signals are injected into the acquisition chain. The latter is the central concern of this work and is the focus of the remainder of this chapter.

Physiological signal generators provide a controlled, reproducible alternative to in vivo recordings for the testing, calibration, and evaluation of biomedical acquisition systems [37]. Unlike head phantoms, which mostly simulate the conductive medium of the body, physiological generators directly inject known electrophysiological signals into a device under test, enabling deterministic evaluation of analog front-ends, digitization fidelity, and downstream signal-processing pipelines without the variability and constraints of human subjects [37], [38]. Existing solutions can be grouped into three broad categories: commercial patient simulators, academic hardware signal generators, and software-only signal generators, summarized in Table 1.2 and discussed in detail in the following subsections. Each category addresses a subset of the validation challenge but presents specific limitations when considered against the requirements of full PSG, which simultaneously involves EEG, EOG, EMG, and ECG signals.

3.2.1 Commercial patient simulators

Commercial patient simulators represent the industry standard for biomedical equipment validation, regulatory testing, and hospital preventive maintenance. The most widely deployed device of this class is the Fluke Biomedical ProSim 8, which integrates ECG (including fetal and arrhythmia waveforms), respiration, invasive and non-invasive blood pressure, cardiac output, temperature, and SpO₂ simulation in a single portable unit [39]. Comparable products include the SECG 5.0 simulator from Whaleteq and the Datrend Vital Sign series [40], [41]. These devices are designed for the latest standards and offer well-characterized waveform libraries, making them suitable for regulatory acceptance and routine clinical equipment maintenance. The Whaleteq simulators also offer the possibility to play back raw data from public database [40]. However, despite their breadth across cardiopulmonary parameters, commercial patient simulators in this class do not include simultaneous multi channel simulation capabilities, which excludes them from comprehensive PSG device validation. Whaleteq also offers an EEG signal generator, the SEEG 100 designed for the latest EEG standard. It offers 9 channel ultra low voltage signal output and the possibility to playback raw data [42]. Another commercial device, not intended to be a patient simulator, is LabSim from Brain Products. It is a dedicated commercial troubleshooting tool that generates reliable test signals to verify EEG amplifier integrity and signal pathway performance, supporting major Brain Products amplifier families [43]. While effective for amplifier-level diagnostics, LabSim is restricted to a single manufacturer’s hardware. Furthermore, the devices cited above have other limitations, including fixed amplitude ranges optimized for ECG or EEG scale signals and costs of several thousand euros that restrict accessibility for academic research environments.

3.2.2 Academic hardware generators

The most active research direction has been the development of academic hardware generators that aim to combine flexibility with affordability. Powell et al. [37] introduced NeuroDAC, an open-source biosignal waveform generator built around an off-the-shelf eight channel audio DAC, with a custom analog attenuation and buffering stage to scale outputs to physiological microvolt-level amplitudes. NeuroDAC achieves a noise floor of 5.35 μV RMS and a high-pass -3 dB cutoff at 0.57 Hz, and is capable of replaying EMG, EEG, ECG, and extracellular neural recordings with correlation coefficients ranging from 0.83 to 0.98 against original recordings. Earlier, Haci et al. [38] proposed a 32-channel ultra-low-noise arbitrary signal generation platform built around a Raspberry Pi 3, a custom FPGA board, and dedicated DACs, providing a higher channel count and stronger isolation between digital and analog domains, but at the cost of significant custom circuitry that limits replicability. Karakulak [44] proposed a lower-cost alternative, in which experimentally recorded EEG waveforms stored on an SD card are reproduced through the internal DAC and PWM peripherals of an STM32 device, providing an affordable, education-oriented platform. Bashirov et al. [45] developed an automated 8-channel phantom EEG simulator intended to be used as an automated test bench, research installation, or precision generator of ultra low voltage signals for EEG analog circuit development, demonstrating signal correlation coefficients of 0.8–0.9 against original recordings. While these academic platforms have demonstrated that low-cost physiological playback is technically feasible, They are focused on raw data playback and do not introduce independent control over artifacts in the data. Artifact injection is a key design feature of the physiological signal generator developed in this work and will be discussed in greater detail in Section 1.4. Moreover, none of them allows the connection of dry electrodes directly to the generator outputs.

3.2.3 Software-only biosignal generators

A complementary class of solutions consists of purely software-based biosignal generators that produce digital waveforms for offline algorithm testing. SEED-G, developed by Anzolin et al. [46], is an open-source MATLAB toolbox capable of generating personalized pseudo-EEG datasets with imposed connectivity patterns, providing a ground-truth benchmark for connectivity-estimation algorithms. Another MATLAB-based toolbox is SEREEGA, presented by Krol et al. [47]. It simulates different types of EEG epochs and combines the majority of existing simulation approaches reported in the literature into a single toolbox. Marsousi et al. [48] proposed a flexible cubic B-spline interpolator-based framework capable of simulating multiple physiological signal types, and notably included a Simulated Artifact Generator (SAG) module as part of the scheme, one of the few generators in the literature to explicitly address artifact synthesis. Karhinoja et al. [49] more recently introduced a flexible framework for generating synthetic ECG and PPG signals, incorporating physiological effects such as breathing modulation and changes in heart rate, along with the addition of noise and artifacts, for data augmentation and increased variation in machine-learning applications. Petersen et al. [50] proposed a procedural EEG simulator implemented as a mobile application capable of streaming synthesized signals into anesthesia depth-of-hypnosis monitors, enabling systematic stimulus-response characterization of clinical devices. While these solutions offer maximal flexibility and no hardware cost, they cannot validate the analog acquisition chain (electrodes, leads, instrumentation amplifiers, and analog-to-digital conversion), which is a critical requirement for hardware-level PSG device validation.

Table 1.2: Comparison of physiological signal generator approaches for PSG device test and validation. ch. = channel; resp. = respiration; HW = hardware; N/R = not reported.

Generator	Modalities	Cost	Key trade-off	Ref
Commercial patient simulators	ECG, resp., BP, SpO ₂ , temp.	High	Industry standard but no EEG/EOG/EMG	[39, 40, 41]
Whaleteq SEEG 100 (commercial)	EEG (9 ch.)	N/R	Standards-compliant but EEG only	[42]
NeuroDAC	EEG, EMG, ECG (8 ch.)	Medium	Audio-DAC playback and low noise but no artifact control	[37]
Haci et al.	Multiple (32 ch.)	N/R	High channel count but complex custom HW and no artifact control	[38]
Karakulak	EEG (1 ch.)	Very low	Affordable and education-oriented but only one channel and no artifact control	[44]
Bashirov et Balovnev	EEG (8 ch.)	N/R	EEG only	[45]
Software EEG simulators	EEG	Free	Flexible but cannot validate hardware	[46, 47]
Marsousi	Multiple	Free	Includes artifact synthesis but software only	[48]
Karhinoja	ECG, PPG	Free	Noise and artifact modeling but software only	[49]
Petersen app	EEG	N/R	Drives anesthesia monitors but no front end validation and EEG only	[50]

3.3 Artifacts injection in physiological signal generators

Among the three categories of physiological signal generators reviewed above, none provides a flexible and controllable mechanism for artifact injection integrated into a hardware device. A real EEG recording, as discussed in section 2, is constantly contaminated by both internal and external artifacts: the subject’s cardiac, ocular, and muscular activity, picked up by all electrodes in ranges from near-field to far-field depending on the channel, as well as electrical grid interference and disturbances due to movement or contact introduced by the measurement chain itself. All of these sources are non-stationary over the course of an eight-hour recording, and together they constitute the primary source of degradation against which any acquisition or processing system must be evaluated.

Validation of PSG monitoring devices is commonly performed by injecting clean reference waveforms into the acquisition chain, which is sufficient to verify gain, linearity, and basic signal integrity. Replaying previously recorded clinical signals adds a layer of realism, since the device is then exposed to the artifacts naturally present in those recordings [37, 38]. The limitation of this second approach is that the user inherits whatever artifact content the recording happens to contain: the type, amplitude, timing, and combination of disturbances are fixed by the dataset and cannot be varied independently. A small number of model-based software frameworks have explicitly addressed artifact synthesis, notably the SAG module of Marsousi et al. [48] and the noise and artifact generation of Karhinoja et al. [49], but as purely digital tools, they cannot validate the analog acquisition chain where many artifacts physically originate. The generator presented in this thesis is built around replaying real recordings, but extends the approach by allowing the principal classes of internal artifacts to be added, scaled, or removed on top of the underlying physiology, with parameters set by the user. This way, the acquisition chain and the downstream processing pipeline can be probed not only against the artifact mix of a given night; but against controlled, reproducible variations of it.

3.4 Synthesis

Across the three generator categories reviewed above, and the additional requirement of artifact injection, no existing physiological signal generator simultaneously satisfies the needs of accessible, multi-modal, artifact-aware, hardware-level capable of driving dry electrodes. Commercial patient simulators provide regulatory-grade reliability but omit the EEG, EOG, and EMG signals at the heart of PSG [39]. Specialized commercial EEG generators remain single-modality [42] and vendor locked [43]. Academic platforms have demonstrated low-cost, multimodal playback but do not include configurable artifact modules and cannot drive dry electrodes [37, 44]. Software-only generators, while flexible, do not provide a characterized analog signal source for validating the acquisition front-end [46, 50], and the few frameworks offering controllable artifact synthesis do so only in software [48, 49]. **This combination of constraints defines a clear gap: an accessible, multi-modal hardware physiological signal generator capable of driving dry electrodes and simultaneously producing realistic EEG, EOG, EMG, and ECG signals together with a controllable artifact injection module, suitable for the validation of PSG acquisition systems.** Addressing this gap is the objective of the present work.

Chapter 2

Design of the physiological signal generator

To achieve this physiological signal generator, the goal is to extract a clean signal and isolate the artifacts from a real raw PSG recording dataset for playback on an STM32 microcontroller driving an analog electronic circuit. The data are stored on an SD card to allow for three main playback modes:

1. Replay of the raw signal.
2. Replay of the clean signal with time-aligned artifact mixing.
3. Replay of the clean signal with controllable artifact injection.

All three modes share the option of selecting a specific sleep stage to replay.

This chapter is divided into three sections describing the methods used for signal processing, hardware development, and firmware implementation, respectively.

1 Signal processing

The physiological signals chosen for the playback originate from the Haaglanden Medisch Centrum (HMC) sleep staging database on PhysioNet [51, 52, 53]. The database contains 151 whole-night PSG recordings of randomly selected patients in a heterogeneous population that underwent PSG examination. The dataset is composed of eight channels including four EEGs, with the following derivations : F4/M1, C4/M1, O2/M1, and C3/M2 (Figure 1.1), two EOGs with E1/M2 and E2/M2 derivations (Figure 1.2), one chin EMG (Figure 1.3), and one ECG with lead II derivation (Figure 1.4a), as well as a text file with annotations corresponding to the sleep stages scored by sleep technicians. This setup meets the minimal recommended technical specifications for sleep stage scoring of the AASM[5]. All signals are sampled at 256 Hz and no additional digital filtering was applied besides the default analog pre-filter settings mentioned in the header files. The database contains raw signals still containing artifacts. In the scope of this project, the artifacts are intended to be modulable. In this perspective, the signals need to be made artifact-free in order to reinject them in a flexible manner. For this work, all the signal processing was done in MATLAB using tools such as the Signal Processing Toolbox. The complete recordings are selected from the database and used in the processing pipeline.

1.1 Data preprocessing

Two categories of artifacts need to be distinguished: external and internal artifacts. External artifacts refer to all sources of signal disturbance external to the human body and not induced by the patient while internal artifacts refer to signal disturbances induced by the patient. More details about artifacts are given in section 2. External artifacts are usually straightforward to handle as they are well characterized and confined to specific frequency bands. Internal artifacts, on the other hand, are more complex to clean and require specific methods.

The first two steps of the processing pipeline consist of bandpass filtering the signals and removing the power line noise. The bandpass filtering is actually already done at signal acquisition and is referenced in the metadata of each recording. The Low and High pass cut off values follow the recommendations described in Section 1.

For the power line noise removal, a notch filter is applied to remove the 50 Hz noise present in the recordings, as the recordings were done in Europe. The notch filter is implemented using a second-order bandstop infinite impulse response filter with half-power frequencies of 49 Hz and 51 Hz, by using MATLAB's `designfilt` function from the Signal Processing Toolbox. The filter was applied only to the channels with 50 Hz included in their bandpass, i.e., the ECG and EMG channels.

1.2 Internal artifacts removal

The strategy for this artifact removal pipeline was chosen to be a method-per-artifact approach. For each artifact, the goal is to extract the artifact from the sources it contaminates and save the artifact component in a dedicated file. First, it is important to remember what the different artifacts are and which sources they contaminate (Section 2.2). The ECG artifact contaminates all the other signals. The EOG introduces artifacts mainly in the EEG signals. EEG suffers from EMG artifacts, but not only from chin EMG. EEG does not induce artifacts but is heavily polluted by artifacts from other signals. The remainder of this section develops the method used for each artifact type.

1.2.1 ECG artifacts removal

ECG artifacts are handled based on a template subtraction method. This method is composed of three main steps and is summarized in Figure 2.1. The method is inspired by the ensemble average subtraction algorithm used in [54] and was adapted empirically.

The first step consists of detecting the R-peaks (see Figure 1.4b) in the ECG signal. First, a local threshold is determined for each 30-second window of the signal. This threshold is set to 0.5 times the 99.5th percentile of the signal in this window in order to ignore spurious peak values while remaining high enough to avoid detecting T-waves. The peaks are then detected using MATLAB's `findpeaks` function with a minimum peak height condition defined by the local threshold and a minimum distance condition of 600 ms between consecutive peaks, corresponding to a maximum heart rate of 100 bpm. This refractory period reflects the physiological constraint that an R-peak cannot occur twice within such a short interval.

The second step consists in the ECG template construction for each channel. In the contaminated signals, a window is built around each peak location (aligned with the peak location from the ECG channel), from -200 ms to 400 ms, in order to capture the complete PQRST segment; the PR segment is shorter than the T-wave segment, which explains the

asymmetry around the peak. During sleep, the average heart rate is usually around 50 bpm for a healthy patient, but it can increase during certain periods of the night. A trade-off must be chosen in order to capture the complete PQRST segment at low heart rates and not exceed one cycle period at higher rates. A window of 600 ms is a good compromise, as it captures a full PQRST segment at the average heart rate and does not overlap with the next cycle until a heart rate of 100 bpm. The template is the average of all the windows constructed in the signal. Template subtraction is a technique validated for ECG-artifact removal from the trunk EMG [55] and has also been used for removal from the EEG [54], and can be extrapolated to the EOG and chin EMG channels, since the assumption underlying its use for the EEG is equally valid for the EOG and chin EMG.

Finally, the template subtraction is applied to each channel. The template is specific to each channel, as the cardiac field effect propagates differently depending on the electrode location. The template subtraction is applied to the four EEGs, the two EOGs, and the EMG channel. The artifacts are stored separately for each channel, so that they can be reused in the generator.

This method assumes the stationarity of the ECG artifact in amplitude and morphology, and is therefore better suited for healthy patients. The method was tested on the *SN002* patient from the database and results are analyzed in Section 1.2.4. The ECG artifact removal is not optimal for patients with high beat-to-beat variability, since a constant template is subtracted from each peak.




Figure 2.1: ECG artifact removal pipeline

1.2.2 EOG artifacts removal

EOG artifact removal from the EEG can be handled by several methods including regression, principal component analysis, and independent component analysis (ICA). For this work, the regression method was chosen for its simplicity and interpretability with a low number of EEG channels [56]. Specifically, a linear regression model was used based on the simpler REG-RAW procedure presented by Wallstrom et al. [56]. The model states that each recorded EEG channel is a linear combination of the EOG channels plus the true EEG signal. In the REG-RAW model, the EOG channels are not filtered. The EOG free estimated channel is given in Equation 2.1.

$$cleanEEG_i(t) = recEEG_i(t) - \alpha_1 HEOG_i(t) - \alpha_2 VEOG_i(t) \quad (2.1)$$

In this work, the α coefficients are estimated by computing the least-squares solution for each i pair. It is important to note that the EOG channels present in the database are not dedicated horizontal and vertical channels (E1-M2 and E2-M2). Therefore, separate coefficients cannot be estimated for the horizontal and vertical derivations. Here, the horizontal

and vertical contributions are mixed in the two coefficients, which is mathematically valid but may result in poorer artifact rejection for certain EOG patterns. This step takes place after the ECG template subtraction, so that the ECG residuals present in the channels are not absorbed by the linear regression and do not contaminate the output.

This method relies on the assumption that the EOG channels carry no cerebral activity. In practice, EEG sources (notably frontal sources, closer to the EOG electrodes) leak into the EOG derivations, so the regression subtracts a fraction of genuine brain activity along with the ocular artifact [56]. This limitation was accepted in exchange for the simplicity of the method and its ability to output the estimated artifact as a separate signal stream, as required by the simulator.

1.2.3 EMG artifacts removal

The EMG contamination of EEG cannot be addressed in this pipeline in the same way as the cardiac and ocular artifacts. The standard approach in the literature is blind source separation, most often ICA. It returns at most as many components as input channels, and its decomposition is predicated on the assumption that each channel is a linear mixture of spatially fixed sources that the algorithm can unmix statistically [57]. With 4 EEG channels, at most 4 independent components can be extracted from the HMC data, fewer than the number of physiologically active sources contributing to the recording: cortical activity, ocular movements, cardiac field potential, and myogenic activity from at least the frontalis, temporalis, and occipitalis muscles [27]. The decomposition is therefore mathematically underdetermined for the present problem: the 4 components returned cannot correspond one-to-one to the sources, and each component is necessarily a mixture of several physiological generators. In practice, reliable separation of specific artifact from neural sources in EEG has only been demonstrated on high-density montages, for instance with Gwin et al. [58] using 248 electrodes to isolate gait-related movement artifacts during locomotion. Methods exist to deal with fewer independent components than sources, but they are beyond the scope of this work.

An alternative would be to use the recorded chin EMG channel as a reference in a regression scheme analogous to the EOG removal described Section 1.2.2. As established in Section 2.2, however, the chin channel records the submental muscles, whereas most of the EEG contamination originates from the muscles adjacent to the scalp electrodes [27]. The two signals are only linked through sleep stages and arousal level, not the sample-by-sample correspondence on which a regression-based subtraction depends. Using the chin channel as a regressor would therefore not subtract the EEG contamination but would inject a scaled copy of an unrelated signal into the EEG output.

The pipeline consequently does not include a residual EMG removal stage. This decision propagates to the generator architecture, unlike the cardiac and ocular artifacts, for which a clean artifact stream is extracted during processing and made available for parametric re-injection by the microcontroller, no separate EMG stream is produced. The EMG contamination present in the raw HMC recordings remains embedded in the "cleaned" EEG output of the pipeline, at its original amplitude and timing. This limitation is acceptable in the context of this thesis. Users requiring an explicitly tunable EMG artifact would need either a montage with frontal or scalp electrodes, or a synthetic muscle-artifact model, both outside the scope of this work.

1.2.4 Validation metrics

The ECG template subtraction can be evaluated by examining the artifact template for each channel. Over all the beats, the parts of the signal surrounding the R-peak average out, leaving only the ECG artifact, since it is constant across the peaks. The template is then subtracted from the averaged ECG artifact. The result is shown in Figure 2.2. The template clearly represents the PQRST waveform of the ECG in the F4, C4, and O2 channels of the EEG, as well as in the EOG channels. The templates for the C3 and EMG channels are smaller in amplitude than the others, and the PQRST waveform is less distinguishable; nevertheless, the template subtraction is not harmful, since the signal being subtracted is small.




Figure 2.2: ECG artifact template for each channel.

The result of artifact removal pipeline on the EEG F4-M1 is shown in Figure 2.3. The ECG artifact stream is composed of the ECG template located around each R-peak. The EOG artifact appearing in the EEG is a linear combination of the two EOG channels, as the regression model assumes. As discussed in the EOG-removal method, the trade-off is that a portion of the useful EEG signal, specifically the frontal brain activity also captured by the EOG electrodes, is subtracted along with the artifact. The ECG artifacts are well removed, as can be seen from the small peaks that have been removed from the raw signal.




Figure 2.3: Artifact removal pipeline for the EEG1 channel. The upper graph is the raw signal. The second and third one are the ECG and EOG artifact streams. The bottom one is the cleaned signal, with the artifact templates subtracted from the raw signal.

1.3 Data export

Once the removal of the artifacts is complete, the signals must be written to a format that the STM32 can stream directly from the SD card. Four binary files are produced per recording: the eight-channel raw signal, the eight-channel cleaned signal, and two artifact streams (ECG on seven channels, EOG on the four EEG channels).

The raw and cleaned signals are stored as unsigned 12-bit codes centered around the DAC mid-scale, so a zero value sample in the original signal maps to the DC operating point $V_{\text{ref}} = 1.65 \text{ V}$ at the DAC output. Artifact streams are stored as signed 16-bit values centered around zero, because they are summed with the cleaned signal, to which mid-scale offset is applied, in the firmware. The int16 container leaves headroom for the mixing coefficient α to scale the artifact up by more than one without overflow before the final saturation to $[0, 4095]$.

The per-channel amplitude scaling that maps microvolts to DAC codes is set from the 99.5 percentile of the absolute value of each cleaned channel. Mapping the maximum to the full DAC swing would allocate most of the 12-bit range to sporadic large spikes, typically electrode pops or motion artifacts, and compress the physiological signal of interest into a fraction of the codes. The cost of this choice is that the 0.5% of samples lying above the percentile are saturated at the DAC rails. The same scaling factor is then reused for the corresponding artifact stream, so that an artifact mixed with $\alpha = 1$ has the amplitude it actually had in the original recording. The percentile value was chosen assuming that the 0.5 % of the signal saturated at the DAC does not represent a significant loss of information.

1.4 Artifacts injection

Based on the mode chosen for the generator, one can choose to mix or inject artifacts into the cleaned signal. These artifacts are composed of those obtained during the cleaning

pipeline. Because the per-channel scaling described above normalizes the artifact to the cleaned channel normalized amplitude, artifact spikes larger than the signal amplitude cannot be reproduced, which is not representative of large EOG artifacts but is a known limitation. The signal integrity is therefore preserved, and the only way to introduce "out-of-range" artifacts is through those not handled by software, i.e., motion artifacts, electromagnetic disturbances, etc.

The files are stored on an SD card for two reasons: the size of the files and the flexibility offered to the user. An SD card allows overnight recordings to be stored, enabling the user to navigate freely within the recording without loading selected samples one at a time. It is also straightforward for the user to load new recordings onto the SD card and play them via the command interface, without having to modify a single line of code.

1.4.1 Artifacts mixing

In the mode where artifacts are mixed with the cleaned signal, one can choose the mixing coefficient for the physiological artifacts. For the internal artifact injection, the process is straightforward, since these artifacts are removed using template subtraction or regression methods, the artifact-only stream is readily available. It is stored in a dedicated file and simply summed with the signal, weighted by a mixing coefficient, in the STM32 firmware. As a result, the firmware only needs three files: the clean signal, one ECG artifact file, and one EOG artifact file.

1.4.2 Artifacts library

The artifact removal pipeline also enables the creation of an artifact library for the third mode of the generator. From the different artifact signals, it is possible to extract several distinct artifact patterns that the user can inject into the cleaned signal in order to observe specific artifact patterns that are not time-correlated with the original signal. It was chosen to use four distinct artifact templates: a single QRS template, a 30-second ECG template, a single large blink, and a 10 seconds long EOG blink sequence. All of them are derived from the artifact removal pipeline. The single QRS template is the ECG artifact template obtained in the ECG removal pipeline. The 30-second ECG template is a 30-second segment of the ECG artifact signal. The EOG blink template is created by finding peaks, above a certain threshold, in the EOG artifact signal and extracting a window around the first selected peak; it is assumed that large peaks in the EOG signal correspond to eye blinks. The EOG blink sequence is selected as the 10-second segment with the highest energy in the EOG signal. The four artifacts patterns are shown in Figure 2.4.




Figure 2.4: Artifact library derived from the EEG1 channel artifact streams.

2 Hardware

This section describes the design of the generator circuit, whose role is to replay physiological signals at their physiological amplitudes. The section is divided into two parts: the first develops the theoretical circuit design, and the second covers the printed circuit board (PCB) design. The primary design constraint is noise. Since the target output amplitudes reach down to the tens of microvolts, every block in the signal path must contribute as little noise as possible. To keep the engineering effort focused on the parts of the system specific to this work, and given the limited time available, previously developed electronic modules were reused wherever possible. The circuit is organized into a digital part and an analog part. The digital side handles signal storage and channel control: an SDMMC interface reads samples from an SD card, an I²C bus carries those samples to the DAC, and an SPI daisy chain configures the controllable output stages of each channel. The analog side carries the DAC output through the signal-conditioning chain that brings it down to the target amplitude at the output pad. The detailed topology of each part is developed in Section 2.1.

2.1 Circuit design

The hardware is split between two boards, and an overview of the hardware architecture is described in Figure 2.5. The Margherita board, a Microsys module reused for this work, hosts the Nucleo board h745zi-q, the SD card interface, I²C and SPI ready to use buses, and a 3.3 V rail. Reusing this module concentrated the engineering effort of the present work on the analog signal path and the channel-control electronics, which were designed from scratch on a dedicated generator board. The generator board carries an 8-channel Digital to Analog Converter (DAC), a 1.65 V analog reference for bipolar signal swing around the supply midpoint, and the analog signal paths, for which seven are identical and one is adapted for a differential output (ECG output). Each path consists of an RC filter, an input buffer, a two-stage controllable resistive divider, and an output buffer driving a conductive pad. The two divider stages are independently controllable from the microcontroller over a single SPI daisy chain: the first stage selects between two coarse attenuation settings via an ADG714 octal switches, while the second stage provides a finer adjustment via an AD5161 digital potentiometer used as a rheostat. This control allows each channel to be mapped to the amplitude range of the modality it represents: EEG, EOG, EMG, or ECG, without changing the underlying topology. The choice to power the system from a wall-derived supply rather than from batteries is a known compromise; it makes the device more practical to operate, since a supply voltage designed for the Margherita is available, but it introduces more noise than a battery would. The impact of this trade-off on the generator output noise is quantified in Section 2.1.4 and in Chapter 3. All other design decisions discussed in the remainder of this chapter were driven by the signal-quality requirements derived from Section 1.




Figure 2.5: Hardware architecture of the physiological signal generator.

2.1.1 Power supply and analog reference

Supply rail: The generator board is powered from a single 3.3 V rail. Two wall-derived supply options are available:

- A **Microsys supply** intended to be provided on the Margherita, which converts mains voltage through a transformer followed by a DC-DC stage. This is the intended operational mode of the device, in order to simultaneously supply the 3.3 V, 5 V, and 12 V rails required for the initial use of the Margherita. The transformer provides galvanic isolation from the mains, which helps reduce coupled noise. However, the DC-DC stage can itself be a source of noise, the magnitude of which will be quantified in Section 3.1.
- A **regulated laboratory linear supply**, a Tenma 72-10495 dual-channel bench supply set to 3.3 V, connected at the 3.3 V supply pin. Its rated output voltage ripple is ≤ 2 mV rms over 20 Hz–20 MHz.

The choice of a wall-derived supply over a battery is a deliberate trade-off. A battery would yield a lower noise floor by removing line-frequency interference and switching-converter residue from the supply rail. The wall path was retained because it is the operating mode supported by the existing Margherita infrastructure, because it avoids runtime power management during long PSG replays, and because the magnitude of the resulting noise penalty is itself a useful experimental quantity for a generator-characterization study. The 5 V and 12 V rails available on the Margherita are not used by the generator board.

Analog voltage reference: The 3.3 V rail is single-ended, asymmetric with respect to ground, but the physiological signals to be generated are bipolar around zero. To represent a bipolar swing on a unipolar supply, a 1.65 V analog reference V_{ref} is generated on the board and used as the common reference at the bottom of every output divider. A signal coded around the DAC mid-scale should therefore appear at the output pad as a swing around V_{ref} , which is the analog equivalent of a signal swing around zero.

V_{ref} is produced by dividing the 3.3 V rail with two 10 k Ω 0.1% tolerance resistors, decoupling the reference with a 100 nF/10 μ F capacitor pair, and buffering the resulting node through an op-amp configured as a unity-gain follower (Figure 2.6). One should note that

resistance mismatch in the divider is not a concern, as the divider’s purpose is to set an analog reference; the exact value of the reference itself is not critical. Beyond the decoupling function of the capacitor pair, the resistive divider together with the $10\,\mu\text{F}$ capacitor forms a low-pass filter that smooths the rail voltage noise; its characteristics are explored in Section 2.1.3.

The buffer that follows is essential for two reasons. First, V_{ref} serves as the reference at the foot of every output divider (sixteen resistors in total), so any disturbance on V_{ref} propagates to all eight channels. Second, V_{ref} is brought out to a dedicated conductive pad to which a dry electrode can be connected, in the same way as the reference electrode of a real recording montage. Any load connected to this external pad would modify the divider ratio and consequently shift V_{ref} . The buffer addresses both concerns through its ideally infinite input resistance, which isolates the divider from any load downstream and presents a low-impedance source to the divider bottoms.

The op-amp chosen for the reference buffer is the ADA4528-1, a chopper-stabilized, zero-drift precision amplifier which characteristics will be further developed in Section 2.1.3.




Figure 2.6: Analog voltage reference with RC filter and buffered output.

2.1.2 Digital part

The digital part is responsible for three things: sending sampled signal data to the DAC at the right rate, controlling the channel attenuation settings on the analog signal path, and managing the storage from which signal recordings are replayed. These three tasks map onto three independent buses: I²C, SPI, and SDMMC, all handled by the STM32H745ZI-Q microcontroller on the Margherita board. The user interface is exposed over USB to a host PC and is described in Section 3.

Microcontroller: The microcontroller is an STM32H745ZI-Q mounted on a Nucleo board. From the peripherals available on this device, the present design uses one I²C interface for DAC data communication, one SPI interface for attenuation control, one SDMMC peripheral for signal storage, one timer providing the 256 Hz sample tick, one UART exposed as ST-LINK for the user interface, and a single GPIO line for the SPI chip select. The peripherals were chosen from those already implemented on the Margherita and available to the user¹. The H745 microcontroller was chosen because it had already been validated on the Margherita board; its 400 MHz core and DMA capability exceed the requirements of this project but provide ample margin for this work and possible future firmware extensions.

¹I²C1 was selected because the bus is already prepared with the required pull-ups, and both I²C1 and SPI6 are externally accessible through connectors on the Margherita.

DAC7578: The DAC7578 is a 12-bit, 8-channel DAC with an I²C interface and an external reference input. Its supply range, of 2.7 V to 5.5 V, is compatible with the 3.3 V rail of the generator.

Two choices justify the use of the DAC7578. First, the integration of eight channels into a single package, with a shared reference and shared I²C bus, makes the per-channel layout symmetric and eliminates the channel-to-channel mismatch that would result from using independent single-channel DACs. Second, 12-bit resolution provides a quantization step at the DAC output of $\frac{V_{\text{REFIN}}}{2^{12}} = \frac{3.3\text{V}}{4096} \approx 806 \mu\text{V}$, which after the largest downstream attenuation considered in this design (Section 2.1.3) reaches largely sub-microvolt scale at the output pad, well below the analog noise floor of the rest of the chain. The limiting factor for output resolution is therefore noise, not the DAC’s quantization.

The reference input V_{REFIN} is tied directly to the 3.3 V supply, configuring the DAC for full-scale output ranging from 0 V to 3.3 V. Signal data is encoded around mid-scale, which aligns with the bipolar centering provided by the analog reference V_{ref} downstream.

I²C communication to the DAC: The DAC is the only device on the I²C bus. The bus operates in fast mode at 400 kHz, with the STM32 internal analog filter enabled and external 4.7 k Ω pull-ups to 3.3 V on the Margherita. The DAC7578 ADDR0 pin is tied to GND. The I²C bus could meet the requirements in normal mode, but fast mode provides additional timing margin; schedulability is discussed in Section 3.1. The LDAC pin is connected to GND because the DAC is used in synchronous mode and The CLR pin to 3.3 V because no hardware reset is used.

SPI gain-control chain: The signal attenuation is selected entirely in the digital domain through two cascaded controllable elements per channel: an analog switch in the first divider stage and a digital potentiometer in the second. Both must be reconfigurable per channel and independently of one another. Although the Margherita’s pin budget exposes several SPI buses, only one was used in order to limit the number of connections between the two boards.

The constraint is resolved by daisy-chaining ten SPI-controlled devices on a single bus, sharing one chip select line:

- Two **ADG714** octal SPST switches in series, providing sixteen independent switch lines, two per channel, used to select between the two R_1 values in the first divider stage.
- Eight **AD5161** 8 bit digital potentiometers, one per channel, used as rheostats in the second divider stage.

Each device shifts an 8-bit register into the next, so a complete chain update consists of 80 bits shifted in MSB-first while the shared CS line is held low. At the SPI clock used in this design, 390 kHz, a chain update takes approximately 200 μs , short compared to the inter-sample period of 3.9 ms at 256 Hz. The gain is not intended to change during playback anyway.

The ADG714 has an on-state resistance of 2.5 Ω typical, which is negligible compared to the kilo-ohm-scale resistors of the divider it controls; the switch therefore introduces no measurable gain error when closed. The AD5161 is a 10 k Ω digital potentiometer with 256 wiper positions, giving a nominal step of 39 Ω ; the wiper itself contributes approximately

60 Ω of series resistance at the minimum-resistance setting. Both devices operate from the 3.3 V rail.

SD card storage: Signal data is stored on a microSD card and accessed by the microcontroller through the SDMMC peripheral in 4-bit mode, with the FatFS filesystem layered on top. SDMMC 4-bit mode is preferred over SPI mode for two reasons: it is the access mode exposed on the Margherita board, and it provides a sustained read throughput of several megabytes per second, which is above the requirement of the application.

The use of a filesystem rather than raw block access in the STM32 memory was a deliberate choice. It allows multiple PSG nights to coexist on the same card, each accessible by filename. It lets the scoring data accompany each recording as a separate file. Moreover, it makes the card human-readable on a PC, which simplifies the workflow of preparing and inspecting signals. The cost is a one-time FAT initialization at boot and a small per-read latency from the filesystem path, both negligible at the data rates involved.

2.1.3 Analog part

Each of the eight DAC outputs feeds an identical analog chain consisting of a buffered low-pass RC filter, a two-stage controllable resistive divider, and an output buffer. The chain is single-ended for the EEG, EOG, and EMG channels (one buffer per channel, output to a single pad) and differential for the ECG channel (one dual op-amp producing two outputs to two pads). The shared reference V_{ref} , described in Section 2.1.1, appears at the foot of every divider stage and also serves as the reference pad of the single ended channels.

The analog chain is the part of the design where the physiological amplitude requirements of each channel (microvolt-scale for EEG, EOG, EMG, and millivolt-scale for ECG) are obtained from the DAC's volt-scale output. This is achieved entirely through the divider, whose two stages are controlled independently per channel over the SPI chain described in Section 2.1.2.

Buffered RC filter: The DAC produces a staircase waveform at 256 Hz, in which each sample is held until the next update. To smooth these steps before they reach the divider, a first-order RC low-pass filter is placed at each DAC output, followed by a unity-gain buffer. Its purpose is to isolate the divider's input impedance, which varies as the user changes the gain settings, from the RC filter, so that the filter's f_c does not shift with the gain selection. The op-amp chosen is the ADA4528-1, the same chopper-stabilized precision amplifier used for the analog reference.

The values were tuned in LTspice to achieve a compromise between two competing constraints. The cutoff must sit above the useful signal band (40 Hz for EEG and EOG, 100 Hz for the wider EMG content) to avoid passband attenuation, and well below the DAC update rate (256 Hz) to give visible smoothing of the inter-sample transitions. The following values were chosen : $R = 10\text{k}\Omega$, $C = 150\text{nF}$ resulting in a cutt off frequency of $f_c = 1/(2\pi RC) \approx 106\text{ Hz}$. The filter introduces approximately 0.5 dB of attenuation at 40 Hz and approximately 8 dB at the DAC update rate. This is enough to smooth the DAC steps without degrading the signal. The RC filter has fixed values, which is not a concern for this work as only PSG signals are intended to be generated, i.e, the frequency band of the signal is not suppose to change.




Figure 2.7: Buffered RC filter at the output of the DAC.

Two-stage divider: The divider implements the channel-specific amplitude scaling. The two stages together cover the orders of magnitude separating the volt-scale DAC output from the microvolt-scale EEG signal at the output pad, while preserving controllability across that range.

Each stage is a resistive voltage divider referenced to V_{ref} , with the controllable element on the series side:

- **Stage 1** allows to select between two attenuation settings. The series resistor is chosen between two fixed values, $R_1^{(a)}$ and $R_1^{(b)}$, which are placed in parallel on switched ADG714 lines. If $R_1^{(a)}$ is selected, the $R_1^{(a)}$ switch is closed and the $R_1^{(b)}$ switch is open. The resistor connecting the divider to V_{ref} is $50\ \Omega$ on all channels.
- **Stage 2** provides fine adjustment. The series resistor is the AD5161 digital potentiometer used as a rheostat, ranging from approximately $60\ \Omega$ (wiper resistance only) to $10\ \text{k}\Omega$ in 256 steps. The resistor connecting the divider to V_{ref} is also $50\ \Omega$ on all channels. As the controllable resistor is placed on the series-side, the control of the output voltage is not linear.

The two stages are cascaded directly, without an inter-stage buffer (Figure 2.8). The second stage therefore loads the first. The divider ratio is given by Equation 2.2.

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{R_2 R_4}{R_1(R_2 + R_3 + R_4) + R_2(R_3 + R_4)} \quad (2.2)$$

with R_1 the stage 1 series resistor (switched), R_2 the stage 1 resistor to V_{ref} ($50\ \Omega$), R_3 the stage 2 series resistor (digipot, from $60\ \Omega$ to $10\ \text{k}\Omega$), and R_4 the stage 2 resistor to V_{ref} ($50\ \Omega$). The output settings are calibrated empirically per channel.

The R_1 values are chosen per modality to map the divider's accessible attenuation range onto the physiological amplitude range of the corresponding signal (Table 2.1).

Table 2.1: R_1 values per channel

Modality	Channels	$R_1^{(a)}$	$R_1^{(b)}$
EEG	4	1 k Ω	10 k Ω
EOG	2	1 k Ω	10 k Ω
EMG	1	1 k Ω	10 k Ω
ECG	1	200 Ω	2 k Ω

The resulting accessible attenuation ranges, computed from Equation 2.2 at the two endpoints of the digipot, are summarized in Table 2.2. Thanks to the selectable R_1 , a wide output range is achievable. The output amplitude ranges from 40.6 μ V to 24.9 mV for the single-ended channels (EEG, EOG, and EMG), and from 398 μ V to 220 mV for the ECG channel on the differential output. Each modality's range comfortably covers the physiological amplitudes reported in Section 1: from a few tens of μ V to more than a hundred μ V for the EEG, few tens of μ V to several mV for the EOG, from a few tens of μ V to several mV for EMG, and approximately 1–2 mV for the ECG R-wave (see Section 2.1.3 for the differential mapping). The ECG channel uses smaller R_1 values than the other modalities because its target amplitude is one order of magnitude larger; less attenuation is therefore required from the divider.

Table 2.2: Per-modality divider attenuation range, computed from Equation 2.2 and expressed in decibels as $A_{dB} = 20 \log_{10}(V_{out}/V_{in})$.

Modality	R_1	Digipot 10 k Ω	Digipot 60 Ω	Output amplitude range (max DAC swing 1.65 V)
EEG/EOG/EMG	1 k Ω	−72.5 dB	−36.4 dB	390 μ V to 24.9 mV
EEG/EOG/EMG	10 k Ω	−92.2 dB	−56.1 dB	40.6 μ V to 2.57 mV
ECG (single-ended)	200 Ω	−60.1 dB	−23.5 dB	1.64 mV to 110 mV
ECG (single-ended)	2 k Ω	−78.4 dB	−42.3 dB	199 μ V to 12.7 mV

The non-linear relationship between digipot setting and output amplitude, visible in Equation 2.2, is not a problem for this application, since the divider settings are calibrated empirically rather than computed in real time. It was chosen to place the digipot as R_3 in order to keep low resistance value for the Thévenin output impedance, limiting the output noise.




Figure 2.8: Schematic of the two stage divider for a single ended channel.

Output buffer (single-ended channels): For the EEG, EOG, and EMG channels (single-ended), an ADA4528-1 unity-gain follower terminates the chain. The output buffer circuit is shown in Figure 2.9. The buffer’s high input impedance isolates the divider output from the load presented by the output pad and the connected electrode, so that variations in that load do not change the attenuation of the divider.

The effective source impedance seen by the buffer is the Thévenin output impedance of the divider, which is dominated by the $50\,\Omega$ shunt resistor R_4 at the output node. Across the entire range of digipot settings, this Thévenin impedance remains bounded between approximately $32\,\Omega$ (digipot at minimum) and $50\,\Omega$ (digipot at maximum): the $50\,\Omega$ shunt to V_{ref} always dominates the parallel combination.

The ADA4528-1’s specifications relevant to this stage are:

- Input voltage noise density: $5.6\,\text{nV}/\sqrt{\text{Hz}}$ at 1 kHz, with no measurable $1/f$ component.
- Input offset voltage: $0.3\,\mu\text{V}$ typical, $2.5\,\mu\text{V}$ maximum.
- Power-supply rejection ratio: 150 dB at 2.5 V supply.
- Rail-to-rail input and output, single-supply operation from 2.2 V to 5.5 V.

The combined effect of these specifications on the output noise floor is discussed in Section 2.1.4.

The chopping mechanism that gives the ADA4528 its low-frequency noise advantage operates at approximately 200 kHz. To attenuate this potential ripple, an RC filter was added inside the feedback loop, based on the recommendations in the ADA4528 datasheet. However, the filter was poorly placed, adding an RC network in the feedback path of a unity-gain buffer does not attenuate the ripple at the output. The closed-loop gain of a non-inverting op amp is given by $A_{\text{CL}} = \frac{A_{\text{OL}}}{1+\beta A_{\text{OL}}}$, with $\beta = \frac{Z_i}{Z_i+Z_f}$, where Z_f is the feedback impedance, Z_i is the impedance from the inverting input to ground, and A_{OL} is the open-loop gain (in theory infinite). In a unity-gain buffer, the inverting input is connected only to the output through the feedback network and has no path to ground, so Z_i is in theory infinite, giving $\beta = 1$ and a closed-loop gain that stays equal to unity regardless of Z_f . The RC filter is therefore transparent in the feedback loop. This misconception was realized too late, but it does not affect the operation of the generator: the chopping ripple at 200 kHz lies well outside the signal bandwidth and does not degrade the signals of interest. To filter the chopping ripple, the filter should have been placed outside the feedback loop, with R in series with the op-amp output and C connected to ground.




Figure 2.9: Output buffer schematic of a single-ended channel, with the RC filter placed as it is in the current design.

Output stage (differential): The ECG output is produced as a differential pair, in line with the lead II configuration, described in Section 1. The generator simulates this by producing two opposite signals around V_{ref} . This design choice could have been made for the single-ended channels as well, but it would have doubled the number of output pads. The true differential setup was therefore used only for the ECG channel, while the single-ended channels share the V_{ref} pad.

The differential pair is generated by a single ADA4528-2, dual op-amp:

- One is a unity-gain follower of the divider output. Its output drives the ECG+ pad.
- The other is a unity-gain inverter referenced to V_{ref} . Its inverting input receives the divider output through a $1\text{ k}\Omega$ resistor, with a matching $1\text{ k}\Omega$ in the feedback, and its non-inverting input is tied to V_{ref} . Its output drives the ECG− pad.

They both integrate the same RC filter feedback loop. It is not effective in the buffer but it is well in the inverting op amp as the inverting input has a finite input impedance.

The differential signal at the recording device is therefore $(V_{\text{ECG}+} - V_{\text{ECG}-}) = (V_{\text{ref}} + V_{\text{ECG}}) - (V_{\text{ref}} + V_{\text{ECG}}) = 2 \times V_{\text{ECG}}$. To produce a 2 mV differential R-wave, the divider needs to produce only 1 mV single-ended at each output.




Figure 2.10: Buffered and inverted outputs of the differential ECG channel, with RC feedback networks placed as they are in the current design.

Output connectors: The output of each signal path drives a conductive pad on the PCB, intended to make contact with a dry electrode. The pads are large enough to fit most of dry electrodes types. Nine pads are provided for the eight signal channels: four for EEG, two for EOG, one for EMG, and two for the differential ECG output. A tenth pad is connected to V_{ref} and serves as the reference electrode of the single-ended channels, in line with the convention described in Section 1. In theory, the V_{ref} at the pad is at the same voltage as the V_{ref} node in the dividers and allows for a clean common-mode voltage rejection when recording the signal in differential setting (One electrode connected to the active pad and one to the reference pad). The results in practice are discussed in Chapter 3.

2.1.4 Noise

This section quantifies the predicted noise floor of the generator's output, building on the topology described above. The goal is to identify which physical mechanisms dominate the noise at the output pad, to confirm that the resulting noise floor is, in theory, below the smallest physiological signals the generator is configured to produce, and to compare these values against the measurements presented in Chapter 3. Because the target physiological amplitudes are on the order of a few tens of μV for the smallest signals, it is important to precisely quantify and analyze the sources of noise in the generator.

Theory: Two physical mechanisms dominate the noise floor of a circuit operating at low frequencies: the thermal noise of the resistive elements, and the intrinsic noise of the active devices.

The thermal noise of a resistor of value R , often called Johnson noise, is a voltage fluctuation produced by the random thermal motion of charge carriers within the resistor. It is white over a very wide bandwidth and has a voltage spectral density given by:

$$e_R = \sqrt{4k_B T R} \quad [\text{V}/\sqrt{\text{Hz}}] \quad (2.3)$$

where $k_B = 1.38 \times 10^{-23}$ J/K is Boltzmann's constant and T is the absolute temperature in kelvin. At room temperature ($T \approx 298$ K), and R is the resistance; Equation 2.3 simplifies to $e_R \approx 0.13\sqrt{R}$ [nV/ $\sqrt{\text{Hz}}$]. A $50\ \Omega$ resistor produces approximately 0.9 [nV/ $\sqrt{\text{Hz}}$], a $1\ \text{k}\Omega$ resistor produces approximately 4.1 [nV/ $\sqrt{\text{Hz}}$], and a $10\ \text{k}\Omega$ resistor produces approximately 12.9 [nV/ $\sqrt{\text{Hz}}$].

An operational amplifier contributes to two intrinsic noise sources at its input: an input voltage noise density $e_{n,\text{op}}$, and an input current noise density $i_{n,\text{op}}$ which produces a noise voltage across the source impedance. The total input-referred noise of an amplifier driven from a source of Thévenin impedance R_S is the root-sum-square of the three uncorrelated contributors:

$$e_{n,\text{in}} = \sqrt{e_{n,\text{op}}^2 + 4k_B T R_S + (i_{n,\text{op}} \cdot R_S)^2} \quad [\text{V}/\sqrt{\text{Hz}}] \quad (2.4)$$

The three terms add in quadrature because they originate from physically uncorrelated processes. For low source impedance, the amplifier's voltage noise dominates; for high source impedance, the current-noise term becomes significant.

For unity-gain follower configurations as used throughout the present design, the input-referred noise of Equation 2.4 appears unchanged at the output. The total noise integrated over a bandwidth of interest is given by Equation 2.5.

$$e_{n,\text{rms}} = e_{n,\text{in}} \cdot \sqrt{\text{Hz}} \quad [\text{V}] \quad (2.5)$$

The input current noise of the ADA4528 ($i_{n,\text{op}} = 0.5\ \text{pA}/\sqrt{\text{Hz}}$) produces a voltage of $(i_n \cdot R_S)^2 \approx 6.25 \times 10^{-4} (\text{nV}/\sqrt{\text{Hz}})^2$ for $R_S = 50\ \Omega$. The current-noise term is therefore dropped in the calculations for the output stage noise below.

The following paragraphs evaluate the contribution of each block in the signal path in turn, before combining them into the per-channel output noise budget. A reminder of the signal path is shown in Figure 2.11 to support the discussion that follows.




Figure 2.11: Simplified electronic schematic of one single ended channel. LTspice schematic.

DAC noise: The DAC7578's intrinsic output noise density is specified in the datasheet. For the DAC operating at 256 Hz update rate, the output voltage noise is approximately $125 \text{ [nV}/\sqrt{\text{Hz}}]$ at full scale at 100 Hz. Therefore, the DAC's contribution at the output pad depends on the divider ratio per channel and ranges from $8.33 \text{ [nV}/\sqrt{\text{Hz}}]$ to $0.015 \text{ [nV}/\sqrt{\text{Hz}}]$ for the ECG output and from $1.89 \text{ [nV}/\sqrt{\text{Hz}}]$ to $0.003 \text{ [nV}/\sqrt{\text{Hz}}]$ for the other single ended outputs, which is below the other noise sources in most divider settings but can be non-negligible with low divider ratios (especially ECG).

Resistors of the divider: The Johnson noise in the divider stage is determined by the Thévenin equivalent resistance of the divider. It is therefore dominated by the 50Ω resistor R_4 that ties the output node to V_{ref} . It is also the source impedance of the output buffer. The worst case contribution (when the digipot is at maximum value) is:

$$e_{R_4} = \sqrt{4k_B T \cdot 50 \Omega} \approx 0.91 \text{ [nV}/\sqrt{\text{Hz}}]$$

Reference buffer (V_{ref}): The V_{ref} buffer is itself an ADA4528-1 follower, but driven from a $5 \text{ k}\Omega$ Thévenin source (the $10 \text{ k}\Omega/10 \text{ k}\Omega$ divider before the buffer). Its input-referred noise is:

$$e_{n,V_{\text{ref}}} = \sqrt{(5.6)^2 + 4k_B T \cdot 5000 + (i_{n,\text{op}} \cdot R_S)^2} \approx \sqrt{31.4 + 82.8 + 6.25} \approx 11 \text{ [nV}/\sqrt{\text{Hz}}]$$

In this case, the $5 \text{ k}\Omega$ source's thermal noise dominates the amplifier's own noise. This noise appears at the bottom of every output divider, as V_{ref} is shared across all channels.

Output buffer (single-ended channels): Applying Equation 2.4 with $R_S \approx 50 \Omega$ and $e_{n,\text{op}} = 5.6 \text{ [nV}/\sqrt{\text{Hz}}]$:

$$e_{n,\text{SE}} = \sqrt{(5.6)^2 + (0.91)^2} \approx 5.67 \text{ [nV}/\sqrt{\text{Hz}}]$$

The amplifier dominates, R_4 's thermal contribution adds approximately 1 % to the total noise. The noise floor of the single-ended output is therefore essentially independent of the divider gain setting.

ECG differential output: The differential ECG output consists of a follower (ECG+) and an inverter (ECG-) producing balanced swings around V_{ref} . The two output noises are uncorrelated (different op-amp of the ADA4528-2, processing physically separated electrical paths) and add in quadrature at the differential output:

$$e_{n,\text{ECG,diff}} = \sqrt{e_{n,\text{ECG+}}^2 + e_{n,\text{ECG-}}^2} = \sqrt{2} \cdot e_{n,\text{SE}} \approx 8.0 \text{ [nV}/\sqrt{\text{Hz}}]$$

This is higher than the single-ended noise by a factor of $\sqrt{2}$, but the differential signal itself is also doubled (Section 2.1.3). The net effect is a $\sqrt{2}$ improvement in signal-to-noise ratio for ECG over what a single-ended chain at the same divider settings would produce.

Combined output noise: Combining the contributions of the previous section in quadrature gives the predicted output noise spectral density of each channel. The dominant terms are the output buffer ($e_{n,SE} = 5.67 \text{ [nV}/\sqrt{\text{Hz}}]$) for single-ended channels, $e_{n,ECG,diff} = 8.0 \text{ [nV}/\sqrt{\text{Hz}}]$ for the ECG differential output) and V_{ref} contribution ($11 \text{ [nV}/\sqrt{\text{Hz}}]$), the DAC noise is low compared to these noise sources when the attenuation is large and the noise originating from the upstream buffer and RC filter is negligible because attenuated in the divider. Results are summarized in Table 2.3.

Table 2.3: Predicted output noise budget per channel, room temperature. Noise RMS value computed for the bandwidth of interest (100 Hz). The "smallest signal" column gives the lower end of generator output amplitude for each modality.

Modality	$e_{n,out}$ [nV/ $\sqrt{\text{Hz}}$]	$e_{n,out}$ RMS [nv]	Smallest signal [μV]
EEG (single-ended)	12.38	123.8	40
EOG (single-ended)	12.38	123.8	40
EMG (single-ended)	12.38	123.8	40
ECG (differential)	13.6	136	400 (diff.)

The output noise spectral density $e_{n,out}$ combines the output-buffer contribution and the worst-case V_{ref} contribution in quadrature: $\sqrt{5.67^2 + 11^2} \approx 12.38 \text{ [nV}/\sqrt{\text{Hz}}]$ for the single-ended channels, and $\sqrt{8.0^2 + 11^2} \approx 13.6 \text{ [nV}/\sqrt{\text{Hz}}]$ for the differential ECG output.

Supply noise contribution: Disturbances on the 3.3 V rail reach the output through three mechanisms of increasing significance.

The first path is the power-supply rejection of each op-amp: with PSRR $\geq 130 \text{ dB}$ for the ADA4528, the rail-to-output coupling is small enough to remain well below the broadband noise floor for any plausible rail ripple. The second path is through the DAC's reference input (V_{REFIN} , tied directly to the rail), which causes rail ripple to scale the DAC output accordingly. This contribution is attenuated by the divider together with the signal, so the signal-to-noise ratio (SNR) with respect to this contribution remains constant along the chain.

The third path, which is the dominant one, is through the analog reference V_{ref} . The reference is derived from the rail through a $10 \text{ k}\Omega/10 \text{ k}\Omega$ divider with $100 \text{ nF}/10 \mu\text{F}$ decoupling at the tap and an ADA4528-1 buffer. The divider passes DC rail voltage at -6 dB ($\times 0.5$); the $10 \mu\text{F}$ capacitor against the $5 \text{ k}\Omega$ Thévenin source impedance of the divider adds a low pass at $f_c = 1/(2\pi \cdot 5 \text{ k}\Omega \cdot 10 \mu\text{F}) \approx 3.2 \text{ Hz}$, attenuating rail ripple above that frequency by 20 dB/decade . The result is that mains-frequency content (50 Hz) reaches V_{ref} attenuated by approximately 30 dB relative to the rail. Once it has reached V_{ref} , however, the residual ripple appears at every output divider's bottom and at the V_{ref} pad itself with no further attenuation.

The supply noise contribution therefore depends almost entirely on the supply mode. The laboratory linear supply (Tenma 72-10495, $\leq 2 \text{ mV rms}$ ripple over 20 Hz – 20 MHz) produces a V_{ref} disturbance of approximately $63 \mu\text{V rms}$ at 50 Hz after the low pass filter. This is the worst case scenario where the supply ripple is concentrated at 50 Hz . However, this supply coupled contribution is concerning for the signal integrity even though the ripple

is concentrated to specific spectral lines. The Microsys supply has no technical datasheet and its noise is characterized in Section 3.1

It is worth stating explicitly what the topology characteristics provide regarding supply ripple rejection. In the ECG configuration, the output is a true differential setup and the supply noise is rejected as common mode voltage by the acquisition system. For the single ended channels, the supply noise coupled through V_{ref} is common to the active output pad and the reference pad. Therefore, if the generator is used with one electrode at the reference and another at the active pad, the supply noise is a common mode voltage and is also rejected thanks to the topology. The only supply coupled noise source is, in theory, the supply noise originating from the DAC.

Noise floor summary: Two observations follow from this noise section.

First, the generator predicted noise floor is dominated by the V_{ref} buffer rather than by the output buffer. The reference buffer is driven from a $5\text{ k}\Omega$ source while the output buffers are driven from a $50\text{ }\Omega$ source, so the reference operates in the resistor-thermal-noise-limited regime while the outputs operate in the amplifier-limited regime. Reducing the V_{ref} divider resistors (for instance to $1\text{ k}\Omega/1\text{ k}\Omega$) would bring the main contribution to the amplifier-limited regime.

Second, the supply noise adds to the intrinsic noise floor of the generator via the DAC only and yields a constant SNR down the chain, as the supply noise in the analog reference cancels out at the acquisition, in theory, due to common-mode rejection. The Tenma supply has a worst-case contribution of 2 mV RMS at the DAC reference, corresponding to an SNR of 64.34 dB , which does not compromise the signal integrity.

Therefore, the generator’s noise floor is theoretically acceptable for the generation of physiological signals.

2.1.5 Simulation results

The goal of this section is to verify that the implemented design matches the theoretical expectations. The simulations were carried out in LTspice.

AC analysis of the filters: The Bode plot of the RC anti-step filter is shown in Figure 2.12a. The simulated response is flat to within -0.58 dB across the useful signal band (DC– 40 Hz), reaches its -3 dB point at $f_c = 106.1\text{ Hz}$, and then decreases at 20 dB/decade as expected from a first-order low-pass network. The attenuation at the DAC update rate of 256 Hz is 8.35 dB . This confirms that the filter performs its intended role of smoothing DAC steps without significantly attenuating signal-band content. The slight phase lag at 40 Hz (approximately -21°) is small enough to leave the temporal morphology of the generated signals unaffected.

The Bode plot of the RC V_{ref} filter is shown in Figure 2.12b. The cutoff frequency is at $f_c = 3.2\text{ Hz}$, and the response decreases at 20 dB/decade above it. The attenuation at 50 Hz is 30 dB . This confirms that the filter performs its intended role of smoothing supply ripple.




Figure 2.12: a) Bode plot of the Post DAC RC filter. b) Bode plot of the V_{ref} RC filter.

AC analysis of the complete analog chain: Rather than simulating the full circuit with all channels in LTspice, only a single channel was simulated. The simulated channel is shown in Figure 2.13. For convenience, a single schematic including both the positive and inverted outputs was used; only the ECG channel uses the inverted output, so the analysis of the positive output remains valid for all other channels.




Figure 2.13: Simulated channel circuit. The two series resistors of the divider were parameterized depending on the simulation.

The full-chain AC response from DAC output to output pad is shown in Figure 2.14 for four representative gain configurations. The two parameters varied are the stage-1 resistor R_1 (switched between $1\text{ k}\Omega$ and $10\text{ k}\Omega$) and the digipot R_3 (set to its minimum and maximum values). Each curve exhibits the expected low-pass character introduced by the RC filter, preceded by a flat passband whose absolute level is set by the divider configuration. The measured passband attenuations are summarized in Table 2.4 alongside the analytical predictions from Section 2.1.3.

Table 2.4: Simulated (LTSpice) vs. predicted passband attenuation per gain configuration for EEG, EOG and EMG channel.

R1	R3	Predicted [dB]	Simulated [dB]	Deviation [dB]
$1\text{ k}\Omega$	$60\text{ }\Omega$	-36.4	-36.7	-0.3
$1\text{ k}\Omega$	$10\text{ k}\Omega$	-72.5	-72.97	-0.47
$10\text{ k}\Omega$	$60\text{ }\Omega$	-56.1	-56.45	-0.35
$10\text{ k}\Omega$	$10\text{ k}\Omega$	-92.2	-92.59	-0.39

The simulated bandpass attenuations agree with the analytical predictions to within 0.5 dB across all four configurations. This confirms that the cascaded divider behaves as predicted across the design’s full attenuation range, and validates the per-modality mapping in Table 2.2.




Figure 2.14: AC analysis of the analog chain.

Sweep through Digipot codes: To verify the output voltage as a function of the digipot setting, and to provide a reference figure for Section 2, the digipot resistance R_3 is swept from approximately 60 Ω to 10 k Ω . The output follows the expected $1/x$ -shaped curve (Figure 2.15), consistent with the placement of R_3 on the series side of the divider.




Figure 2.15: Output voltage in function of digipot values with $R_1=1\text{k}\Omega$

DC sweep simulating sweep through DAC codes: To verify the linearity of the analog chain, and to provide a reference figure for Section 2, V_{in} is swept from 0 to 3.3 V at fixed divider settings ($R_1 = 10\text{k}\Omega$, $R_3 = 60\Omega$). The output varies linearly with V_{in} , centered on V_{ref} , with a slope equal to the divider ratio (Figure 2.16).




Figure 2.16: Output voltage in function of V_{in} with $R_1=10\text{k}\Omega$ and $R_3=60\Omega$

Simulated signal output: The response of the analog chain to a realistic input signal can be simulated by providing a piecewise-linear (PWL) file containing a physiological signal sampled at 256 Hz, representing the DAC output. These simulations confirm that the circuit behaves as predicted under a realistic input waveform (Figures 2.17 and 2.18). The simulation only shows the output swinging around V_{ref} and does not include the differential measurement between the channel pad and the V_{ref} pad; this will be verified in practice.




Figure 2.17: Simulated ECG signal reaching the expected physiological values. The ECG+ and ECG- are measured with $R_1 = 2\text{ k}\Omega$ and $R_3 = 2\text{ k}\Omega$.




Figure 2.18: Simulated EEG F4-M1 signal. The EEG1 channel output reaches the expected physiological values with $R_1 = 10\text{ k}\Omega$ and $R_3 = 3\text{ k}\Omega$.

2.2 Printed circuit board design

The generator was implemented directly on a printed circuit board (PCB), without a prior breadboard prototyping stage. For a circuit operating down to tens-of-microvolt signal levels at the output, breadboard prototyping would not have been informative because of the parasitic effects, lead inductance of the loop areas formed by free wires and ground loops. Any noise figure obtained on a breadboard could not have been attributed unambiguously to the circuit rather than to the prototyping setup. It also would have been fastidious because of the components required by the design. The design effort was therefore concentrated on producing a clean PCB layout from the start. The final design of the physiological signal generator is shown in Figure 2.19. The goal of this section is to explain the choices that led to the final PCB design of the generator.




Figure 2.19: Final PCB design.

2.2.1 Layer stack-up

A four layer board is used, organized as follows:

- layer 1: components and signal traces
- layer 2: ground plane
- layer 3: power plane (3.3V)
- layer 4: signal traces

The motivation for going to four layers is to give every digital trace a low-impedance return path on the plane immediately underneath it, in order to limit ground-bounce noise. With the ground plane on layer 2, for any top-layer signal with a frequency above 100 kHz, the return current in the ground plane will track the signal trace directly beneath it. This applies to the SPI traces, which run across the board, and to the I²C traces, even though they are limited to the DAC. This both minimizes the loop area of the digital return currents and removes the possibility of ground loops being formed by the routing itself. Separating the supply (layer 3) onto its own plane, sandwiched between the ground plane and the bottom signal layer, also provides a degree of intrinsic supply decoupling through the plane-to-plane capacitance of the stack.

2.2.2 Components placement

The surface of the board is mainly covered by the ten conductive output pads (three EEG pads and one V_{ref} pad on the top edge, EEG4 and EOG1 on the right edge, and EOG2, EMG, ECG+ and ECG- on the bottom edge). These pads are located along the edges of the board, so that the active circuit is confined to the area between them (Figure 2.19).

As a result, the analog chain is not very compact, each channel is free to spread out toward its destination pad. Analog traces on this board are therefore longer than they would be in a more compact layout including more conventional output connectors. This is acceptable here because the useful signal bandwidth is below 100 Hz so trace inductance and capacitance are very low. However, longer traces present a larger antenna for externally-coupled interferences, and a larger resistance for nodes that carry current. The second point becomes important for the V_{ref} node and is analyzed in Section 2.2.5. The analog chain of one channel is highlighted in Figure 2.20.

Test points are placed at the output of two analog stages for each channel, the RC filter output and the stage-1 divider output, and on the SPI/I²C lines, so that every block of the signal chain can be probed independently during debugging and characterization.




Figure 2.20: EEG1 signal path. The signal flows from 1 to 7: 1) DAC, 2) RC filter, 3) buffer, 4) switch, 5) divider, 6) output buffer, 7) output pad. The grey rectangle highlights the analog reference generator.

2.2.3 Interconnects to the Margherita board

The generator board does not embed its own microcontroller, SD card, or power-supply module; these are provided by the Margherita board (Section 2.1). The two boards are connected through one 2×5 -pin JST header on the left edge of the generator board, with flying cables approximately 20 cm long. These cables carry the 3.3 V supply, the ground return, the I²C bus to the DAC, and the SPI bus to the digipots and switches. They are a likely path for externally coupled interference because of their length and lack of shielding. The digital communication on the I²C and SPI buses was nevertheless validated in practice, confirming that this coupling does not corrupt the digital transactions.




Figure 2.21: Generator connected to the Margherita.

2.2.4 SPI daisy chain routing

The two ADG714 octal switches and the eight AD5161 digipots are chained on a single SPI bus, with the SDO of each device connected to the SDI of the next one, for a total of ten devices on the chain. A pull up resistor is required between each SDO and SDI pin in order to transfer data. The routing follows the shortest path that visits all the devices in sequence. The Chip Select (CS) and clock (SCLK) of the SPI protocol are shared by the ten devices and are thus not required to follow a certain sequence. The chain runs across the central component zone of the board, in the same region as the analog signal traces. This proximity is mitigated by the stack-up. The chain is routed on a signal layer with the ground plane directly underneath, so the return current will follow the chain trace in the ground plane. The remaining concern is crosstalk between the digital and analog traces, but this should not be a problem because the SPI is silent during playback.

In a dynamic configuration, which corresponds to the intended use of the generator, when any channel's DAC code changes, the current it injects into the V_{ref} node changes accordingly. The non-zero impedance of the node converts these current changes into voltage modulations, which all other channels' dividers see as a fluctuating reference. This manifests as inter-channel coupling: one channel's transfer function depends on the codes of all the others.

Both effects share the same root cause described above. The static effect produces channel-dependent offsets at a fixed configuration, while the dynamic effect produces inter-channel coupling whenever the configuration changes. Mitigation strategies consist in reducing the trace impedance and reducing the channel current sources. One mitigation for a future version would be to replace the V_{ref} trace with a dedicated copper pour on a signal layer in the analog region of the board. This would reduce the resistance, attenuating the voltage drops and significantly reducing inter-channel coupling. Another solution would be to replace the 50Ω resistors with higher value ones to reduce currents. A final mitigation would be to use a local buffer at each channel divider rather than a single shared buffer for all loads.




Figure 2.23: View of the 4th PCB layer. The output of the grey rectangle connects to the trace (wide blue trace) that distributes V_{ref} to all the channel dividers, each shown in a different colour. The grey arrow points to the V_{ref} pad.

2.2.6 ECG differential output

The EEG, EOG, and EMG channels share the V_{ref} pad as their reference. The voltage at the pad may slightly differ from the one at the divider bottom, as explained above. The ECG channel is the exception: it is a true differential pair, with the ECG+ and ECG- pads providing two outputs with no V_{ref} dependence on the recording-amplifier side. The V_{ref} contribution is common to the two pads and is therefore rejected as a true common-mode voltage by the differential acquisition system. This is why the ECG channel should be the only one whose output is unaffected by the inter-channel coupling and V_{ref} node modulations discussed above; any disturbance on V_{ref} appears equally on ECG+ and ECG-, and is canceled by the differential acquisition.

2.2.7 Identified design omissions and post fabrication fixes

Several aspects of the board did not match the original design intent. These deviations are documented here for transparency and as a reference for future versions. None of them currently affect the generator’s behavior.

Floating configuration pins. Two pin-strap inputs were left floating on the board and required post-fabrication fixes to make the corresponding devices usable:

- The ADG714 octal switch has a RESET pin that defaults asynchronously to an undefined state when left floating, leaving the device in a non-operational condition. This was fixed on the assembled board by soldering a short wire between the RESET pin and the 3.3 V supply pin on the opposite side of the device, tying RESET high, as can be seen in Figure 2.19.
- The AD5161 digipot has a mode-select pin (selecting between I²C and SPI control), which defaults to I²C when left floating. The design requires SPI mode so that the digipot can be daisy chained with the ADG714s on the SPI bus. This was patched on the assembled board by bridging the mode-select pin to the adjacent GND pin on each AD5161, forcing SPI mode.

Both fixes were applied to the board before the characterization measurements of Chapter 3. They do not affect signal integrity, since the patched pins are static digital inputs; however, these corrections should be incorporated into the layout in a future version.

Missing local decoupling. The ADA4528 datasheet recommends placing a bypass capacitor as close as possible to each supply pin. This was not implemented on the generator board, the ADA4528 amplifiers in the input and output buffer positions have no local decoupling and rely entirely on the 3.3 V plane. This should not compromise signal integrity, given the high PSRR of the op-amp, but adding local decoupling would be preferable. The AD5161 digipots and the ADG714 switches are in the same situation; as long as it does not affect their operation, this is not critical. The DAC7578 does have local decoupling at its V_{ref} input, but it could be improved to low pass filter the supply. This is noted for completeness, the SNR between the signal and the noise from the unattenuated supply ripple at the DAC is already large enough to ensure signal integrity.

3 Firmware

This section is dedicated to explain the firmware at a high level of abstraction. The firmware running on the STM32 microcontroller is written entirely in C and uses Hardware Abstraction Level (HAL) functions to interface with the different peripherals. The microcontroller is a dual-core device but only the Cortex-M7 is used in this work, a single core has enough headroom for the playback pipeline, and using the Cortex-M4 would add complexity without a clear benefit.

When powered on, the generator is silent and waits for a user command, with the dividers initialized to their default settings. The user interacts with the generator through a terminal command interface over USB, using the UART exposed by the ST-LINK port of the Nucleo board. From the menu, the user can type a command to select the playback mode and start generating signals. During playback, depending on the selected mode, the user can edit the mixing coefficients, inject artifacts, stop playback, or reset the generator. A list of the available commands is given in Appendix D.

3.1 Firmware architecture

When the user starts playback, the main playback task is invoked. This task is the same for all three modes; only the data fed to the DAC differs. Its role is to feed signal data to the DAC at the recordings' original sampling rate. It splits into two subtasks running concurrently. The first, executed in the timer interrupt service routine (ISR), sends one 8-channel sample frame to the DAC on every tick from an active buffer. The second, running in the main loop, fetches the next frames from the SD card and writes them into an inactive buffer that the ISR starts reading when the active buffer is empty. Alongside the playback task, the UART interface remains active so that the user can manually stop playback in order to change the mode or the divider setting. The divider ratio of each channel is intended to be constant during playback and can be modified via a user command before playback.

The choice of using two buffers ensures that playback runs smoothly. Indeed, the SD card path goes through FatFs and the SDMMC peripheral, which adds variable latency on the order of milliseconds per sector read. If the data retrieved from the SD card are stored in one buffer, there would be a mismatch when the buffer is empty because it might not have time to refill before the next ISR. A double buffer scheme solves the problem: two equal-sized buffers, A and B, hold pre-mixed DAC frames. The ISR reads sequentially from one of them, one frame per tick, while the main loop refills the other. When the ISR reaches the end of its buffer, the two swap roles, and a flag is raised to tell the main loop that the just-emptied buffer is free to be refilled.

The pseudocode of the main task is shown below.

Task 1 Playback loop

```
1: buffer_A, buffer_B (each  $1024 \times 8$  frames)
2: ready_A, ready_B : flags raised by the ISR
3: playback_done = 0
4: start_TIM6
5: while playback_done != 0 do
6:   if ready_A then
7:     REFILL(buffer_A)                                ▷ read SD, mix, write 1024 frames
8:     ready_A ← FALSE
9:   else if ready_B then
10:    REFILL(buffer_B)
11:    ready_B ← FALSE
12:   end if
13: end while
14: stop_TIM6
```

Task 2 TIM6 ISR

```
1: active ∈ {A, B}, frame_idx ∈ [0, 1023]
2: frame ← active[frame_idx]
3: DAC7578_WRITEFRAME(frame)                        ▷ 8 channels, synchronous update
4: frame_idx ← frame_idx + 1
5: if frame_idx = 1024 then
6:   frame_idx ← 0
7:   if active = A then
8:     active ← B
9:     ready_A ← TRUE                                ▷ playback loop: refill A
10:  else
11:    active ← A
12:    ready_B ← TRUE
13:  end if
14: end if
```

3.2 Schedulability

Three peripheral activities coexist during playback: the TIM6 ISR, the I²C communication, and the SDMMC reads. TIM6 runs at 256 Hz and its only job is to trigger an interrupt that sends a data frame to the 8-channel DAC over I²C. A loop is responsible for fetching the next data samples from the SD card via SDMMC and storing them in a buffer. The UART remains active during playback to receive and process user commands. The SPI peripheral is not used during playback.

It is important to ensure that these executions are schedulable in order to provide smooth playback. The ISR is called every 3.9 ms. Each buffer holds 1024 frames of 8 channels, which at 256 Hz corresponds to 4 seconds of playback. That sets two budgets. First, the main loop has up to 4 seconds to refill an empty buffer before the ISR runs out of pre-mixed frames. In practice, the worst-case observed refill time, when three files are read in parallel, was 25 ms, leaving ample headroom. Second, the ISR itself must complete a full 8-channel DAC update inside one tick period; each channel takes one 4-byte I²C transaction (address

+ command + 2 data bytes), so a frame is 32 bytes plus protocol overhead. At 400 kHz, with 9 clock periods per byte (one clock cycle per bit + one clock cycle for protocol), this is about 720 μ s per frame, well under the 3.91 ms ISR period. The DAC7578 is written using its synchronous update mode; the first seven channels are loaded into the input registers, and the eighth uses the Write-and-Update-All command, which loads the last input register and simultaneously transfers all eight input registers to the DAC registers, so the eight channels are synchronized. A timing diagram is represented in Figure 2.24.

The microcontroller's load factor is given by Equation 2.7. Any UART interrupt therefore fits comfortably within the budget. It should also be noted that the SPI transaction, computed in Section 2.1.2, lasts 200 μ s and fits well within the budget. It is therefore technically feasible to modify the divider ratio during playback.

$$U = \frac{0.72 \text{ ms}}{3.9 \text{ ms}} + \frac{25 \text{ ms}}{4000 \text{ ms}} \approx 0.185 + 0.006 \approx 0.19 \quad (2.7)$$




Figure 2.24: Schedulability diagram. The TIM6 ISR is active within the playback loop; the two are shown separately to improve visualization.

3.3 Playback modes

The three modes share the architecture of Section 3.1 and differ only in what the mixer fills into the buffer. Each mode is launched by a single UART command and runs until the user stops the playback or until the file ends. The pseudocode of the mixer is given for each mode.

3.3.1 Raw signal playback mode

This mode is the simplest to implement, the mixer only copies the data from the raw data file into the buffers.

Task 3 REFILL in raw playback mode

- 1: **Task** REFILL(*buf*)
 - 2: *buf* $\leftarrow$ READFRAMES(*raw_file*)
 - 3: **end Task**
-

3.3.2 Artifacts mixing playback mode

Two additional files are opened alongside the clean recording: a full-night ECG file and a EOG artifact file. At each refill, the mixer reads 1024 frames per file and computes, for every DAC channel $ch \in \{0, 7\}$, a weighted sum:

$$V_{ch} = S_{ch}^{clean} + \alpha_{ch}^{ECG} \cdot A_{ch}^{ECG} + \alpha_{ch}^{EOG} \cdot A_{ch}^{EOG}$$

where the coefficients α_{ch}^X are stored in two 8-element tables, one per artifact family. Each artifact contribution can be enabled, disabled, or rebalanced while playback is running. The range of coefficient values is not bounded, but the resulting signal value is saturated if it exceeds the DAC range.

Task 4 REFILL in artifact mixing mode

```

1: Task REFILL(buf)
2:    $S^{clean} \leftarrow \text{READFRAMES}(\text{clean\_file})$ 
3:    $A^{ECG} \leftarrow \text{READFRAMES}(\text{ecg\_file})$ 
4:    $A^{EOG} \leftarrow \text{READFRAMES}(\text{eog\_file})$ 
5:   for  $i = 0$  to 1023 do
6:     for  $ch = 0$  to 7 do
7:        $buf[i][ch] \leftarrow S[i][ch]^{clean} + \alpha[ch]^{ECG} A[i]^{ECG} + \alpha[ch]^{EOG} A[i]^{EOG}$ 
8:     end for
9:   end for
10: end Task
```

3.3.3 Artifacts injection mode

This mode allows the user to inject defined artifact samples at any moment into the clean signal. The library file, containing the artifact samples described in Section 1.4.2, is opened in parallel with the clean signal file.

An artifact sample is delivered to the mixer through a small circular buffer holding 1024 frames. The playback loop refills it from the SD card on every iteration. The mixer pulls one frame at a time and adds it to the channels to which it applies (for example, it is not possible to inject an ECG artifact into the ECG channel itself). A separate circular buffer is used instead of extending the three-file mixer used in the mixing playback mode. Indeed, the artifact samples are short, and the user can choose to loop the sample, all of which are easier to manage on a small buffer. Also, the artifact sample is read on demand and only when the user injects it, so SD bandwidth is not spent when no injection is happening.

Task 5 REFILL in artifact injection mode

```

1: Task REFILL(buf)
2:    $S^{clean} \leftarrow \text{READFRAMES}(\text{clean\_file})$ 
3:   for  $i = 0$  to 1023 do
4:      $(A^{ECG} | A^{EOG}) \leftarrow \text{PULLFRAME}(\text{injected\_artifact})$ 
5:     for  $ck = 0$  to 7 do
6:        $buf[i][ch] \leftarrow S[i][ch]^{clean} + \alpha[ch]^X A^X$ 
7:     end for
8:   end for
9: end Task
```

3.3.4 Specific sleep stage playback option

All three modes accept an optional sleep-stage argument. Instead of starting at the beginning of the recording, playback can be restricted to a chosen sleep stage. When one stage is selected, the firmware reads the sleep-scoring file associated with the recording, locates the longest contiguous run of epochs in that stage, and seeks each open file (clean, artifact streams) to the corresponding sample offset before starting the ISR. Playback then proceeds normally and stops at the end of the stage. This is a pre-processing step done once per launch, the playback then works as described before.

Chapter 3

Results and discussion

1 Measurement method and protocol

After troubleshooting and confirming that the PCB was performing as intended, and after debugging the firmware and verifying that the generator was working as expected, the time remaining to measure and characterize the generator was limited. Due to the presence of 8 channels and variable divider settings, many configuration setups could be tested. However, the characterization of the generator was not exhaustive for every channel; a protocol was developed for a few selected channels and settings. This section proposes a characterization protocol for the generator in order to validate several of its aspects:

1. **Intrinsic properties of the generator:** output range, output linearity, and intrinsic noise.
2. **Physiological signal playback:** ability to faithfully replay the original signal regardless of divider settings, and validation of the differential setup.
3. **Validation of the generator modes.**

1.1 Characterization protocol

The characterization protocol is divided into two main steps: the DC characterization for the intrinsic properties, and the dynamic characterization for the physiological signal playback and the validation of the generator modes. Each step requires their own measurement devices: a nanovoltmeter for the DC characterization, and both an oscilloscope and a nanovoltmeter for the dynamic characterization. The remainder of this section describes the different protocols, the measurement devices, and the measurement setup.

1.1.1 DC characterization

By design, the AC output signal is centered around V_{ref} , since the single-supply architecture requires a DC reference around which the signal can swing. Therefore, in order to measure the AC output, a setup that cancels the DC offset is required. However, the characterization of the generator output range, its output linearity, and its noise can be performed in DC. This offers the possibility to use highly accurate devices, which is necessary as the output range of the generator spans from a few tens of microvolts to a hundred millivolts.

The measurement device used for this part of the characterization is the **Keysight 34420A** Nanovolt/Microhm Meter. The 34420A is a differential precision voltmeter offering two measurement channels, each having one negative and one positive terminal. It has a minimum range of one millivolt that provides 0.1 nanovolt resolution, which is highly sufficient for this case. Moreover, the integration time, which is the period during which the device samples the input signal for a measurement and is specified in number of power line cycles (NPLC), can be chosen adequately to provide Normal Mode Rejection (NMR). This is a useful feature as it simplifies the measurement setup by removing the need for a Faraday cage. Furthermore, the measurement cables provided with the device are twisted up to the terminals in order to reduce magnetic field effects. The terminals consist of flat copper ends, which are suitable for connecting to the generator's copper output pad. The copper-to-copper connections ensure that thermal EMF errors are minimal, and they are maintained against the pad by a crocodile clip.

The three generator properties of interest each require their own measurement protocol. They are described later in this section, and the results are analyzed in Section 2.

In order to collect data from the 34420A, the GPIB port of the device is used to communicate with a computer, rather than recording the measurements by hand. A Python script is used to set the desired voltmeter parameters depending on the measurement protocol, acquire data from it, and program the generator settings in order to conduct the desired experiment. The data are then analyzed within the same script and the results are plotted; details about the analysis of the measurements are given in Section 2.

Due to the generator topology, a distinction is made between the single-ended channel results and the ECG channel results.

Output range: The output range is measured by performing a sweep through the digipot codes for a fixed R_1 resistance.

On the measurement device side, the integration time is set to 10 NPLCs to achieve a high NMR (95 dB); the range is set to automatic mode and the analog filter (with -3 dB at 11 Hz) is enabled for noise rejection. On the generator side, all DAC channels are coded to FULL scale; R_1 and the digipot codes are identical across all channels. The DAC channel on which the output is measured is coded to the FULL scale, has the same R_1 value as the other channels, and its digipot codes are swept from 255 to 0 with one step every 600 ms in order to allow the value to settle. The digipot code is linked to the R_3 value as described in Equation 3.1. The minimum resistance value corresponds to a code of 255, while the maximum resistance value corresponds to a code of 0.

$$R_3 = R_{AW} = R_W + R_{AB} - R_{BW} = R_W + 10\text{k}\Omega - \frac{\text{code}}{256}10\text{k}\Omega = R_W + \frac{255 - \text{code}}{256}10\text{k}\Omega \quad (3.1)$$

If the measurement is performed on a single-ended channel, the positive terminal of the measurement cable is connected to the channel output pad and the negative terminal is connected to the V_{ref} pad. If the measurement is performed on the ECG channel, the positive terminal of the measurement cable is connected to the ECG+ pad and the negative terminal is connected to the ECG- pad.

Two plots are derived from the measurements: the output voltage as a function of the digipot code, shown on both a logarithmic and a linear vertical scale, with the predicted

output voltage included for comparison. The two plots allow for better visualization of the output range. The schematic of the divider is recalled in Figure 3.1, and the predicted result is given by Equation 3.2.

$$V_{predicted} = (V_{DAC_{FS}} - V_{ref}) \cdot \frac{R_2 \cdot R_4}{R_1 \cdot (R_2 + R_3 + R_4) + R_2 \cdot (R_3 + R_4)} [V] \quad (3.2)$$

with $V_{DAC_{FS}}$ the DAC full-scale output voltage, R_1 depending on the switch states, and $R_3 = R_W + \frac{255-code}{256} \cdot R_{AB}$.




Figure 3.1: Two stage divider for a single ended channel with $R_{1a}=1k\Omega$, $R_{1b}=10k\Omega$, R_3 =Digipot and $R_2=R_4=50\Omega$.

Output linearity: The output linearity is evaluated by performing a sweep through the DAC codes for a selected channel with fixed R_1 and digipot values.

On the measurement device side, the integration time is set to 10 NPLCs to achieve an NMR of 95 dB; the range is set to automatic mode and the analog filter (with -3 dB at 11 Hz) is enabled for noise rejection. Before starting the sweep, the DAC is set to MID scale and the measured output is saved as the NULL reference for the measurement device. On the generator side, all DAC channels are coded to the same scale; R_1 and the digipot codes are identical across all channels. The channel on which the output is measured has the same R_1 value and digipot codes as the other channels. The DAC code is swept from 0 to 4095 in steps of 16. The measurement connections follow the same setup as the output range protocol.

Three plots are derived from the measurements. The raw measurements are plotted as a function of the DAC codes, and a linear regression is applied to the measurements in order to qualitatively compare the result with the ideal linear output. The Integral Non-Linearity (INL) and Differential Non-Linearity (DNL) are computed from the measurements to quantify the deviation from the ideal output and the uniformity between adjacent codes, they are described in Equations 3.3 and 3.4. The ideal output is a linear slope increasing with the DAC codes; the smaller the INL, the more linear the output. INL and DNL are complementary metrics, INL captures the overall deviation of the transfer function from linearity, while DNL captures the uniformity of the individual voltage steps between adjacent codes. For a physiological signal generator, the goal is to achieve low INL and DNL in the output transfer function. Indeed, the smaller the linearity errors, the closer the output signal is to the original signal.

$$\text{INL}(i) = V_{\text{measured}}(i) - V_{\text{ideal}}(i) \text{ [V]} \quad (3.3)$$

$$\text{DNL}(i) = \frac{V_{\text{measured}}(i+1) - V_{\text{measured}}(i)}{\text{LSB}_{\text{sweep}}} - 1 \text{ [LSB}_{\text{sweep}}] \quad (3.4)$$

Where one $\text{LSB}_{\text{sweep}}$ corresponds to the ideal voltage increment between two adjacent measurement points (i.e., 16 DAC codes).

Noise properties: The intrinsic noise properties of the generator can be evaluated by recording a long sequence of measurements while the generator maintains a fixed DAC code, R_1 value, and digipot code. The AC RMS noise value, i.e., the standard deviation, of the generator output is then computed from the recorded sequence.

On the measurement device side, the integration time is set to 1 NPLC to reduce the acquisition time while keeping an NMR of 60 dB; the range is fixed to 10 mV for a resolution of 1 nV; the analog filter is OFF in order to measure noise over a broader band than DC–11 Hz. Before starting the acquisition, the DAC is set to MID scale and the measured output is saved as the NULL reference for the measurement device. On the generator side, all DAC channels are coded to MID scale; R_1 and the digipot codes are identical across all channels. The channel on which the output is measured has the same DAC code, R_1 value, and digipot codes as the other channels. The measurement connections follow the same setup as the previous protocols. The standard deviation is described in Equation 3.5.

$$\sigma = \sqrt{\frac{1}{N-1} \sum_{i=1}^N (x_i - \bar{x})^2} \quad (3.5)$$

where x_i denotes an individual measurement, $\bar{x}$ is the sample mean of the measurement sequence, and N is the total number of measurements. The protocol was conducted with $N=1000$.

1.1.2 Dynamic characterization

As mentioned above, the AC output signal swings around a DC offset V_{ref} . An oscilloscope was first used to measure the signal at different stages of the generator, offering a direct view of what is being generated by the circuit and easy probing at each stage. The DC nanovoltmeter is not designed to measure continuous waveforms, since its sampling rate is limited. However, the signals generated here are low-frequency, and it turns out that the sampling rate of the device is high enough to reconstruct them without aliasing. Both instruments were used to characterize the dynamic behavior of the generator; their respective contributions are detailed in the remainder of this section.

The oscilloscope used for the dynamic characterization of the generator is the **R&S RTB2004**. It has a minimum input sensitivity of 1 mV per division, 10 vertical divisions, and uses a 10-bit ADC, so one quantization step is approximately 10 μV at its minimum sensitivity setting. These properties enable to measure signals with sufficient resolution down to a few hundred microvolts. The high-resolution mode enables very accurate measurements. It should be noted that, because the signals are centered around 1.65 V, it is not possible to observe the signal swinging around 1.65 V with a sensitivity of 1 mV per division, since the oscilloscope has a finite offset range and therefore cannot display a large DC-centered

signal with such a small sensitivity. The solution is to record the signal in AC mode, which removes the DC offset and allows the signal to be observed swinging around 0. It should be noted that the AC mode of the oscilloscope consists of a hardware high-pass filter whose cut-off frequency is not specified in the datasheet. As a result, it can be expected that the low-frequency content of the signal is attenuated, leading to a slight signal distortion. However, this effect was considered negligible for the present case, although it is worth being aware of it.

These properties therefore do not allow the signal to be measured at physiological amplitudes using the oscilloscope for two reasons; the minimum resolution is too large ($10\ \mu\text{V}$ resolution enables measurement of signals of a few hundred microvolts, not a few tens), and the measurement is referenced to ground, not differential with respect to V_{ref} , so the supply noise is not rejected as it would have been as a common-mode voltage in a differential setup. These limitations are acceptable for characterizing the generator’s ability to replay signals at amplitudes higher than the physiological range. The ability to reach the physiological amplitudes described in Section 1 can instead be assessed through the DC characterization and supported by the dynamic characterization using the DC nanovoltmeter.

The DC nanovoltmeter was considered as an alternative to characterize the dynamic behavior of the generator as the oscilloscope cannot measure signals at the lowest scale of the generator output. It was used in the same way as for the DC characterization, the difference being in the integration time chosen for the measurement. In order to reconstruct the signal without loss, the sampling rate must be at least twice the maximum signal frequency. The minimum integration time of the DC voltmeter is 0.02 power line cycles, i.e., $400\ \mu\text{s}$ at the 50 Hz mains frequency. This corresponds to a theoretical sampling frequency of 2500 Hz, but the maximum internal trigger rate is limited to 250 Hz. This is still sufficient to reconstruct the signal without aliasing, since the signal bandwidths are below 125 Hz for every channel. At first, this setup was considered only as an additional measurement option, since a sufficient sampling rate had not yet been achieved. To reach the maximum trigger rate of the nanovoltmeter, several parameters must be configured appropriately: the display and autozero must be turned off, and the range must be fixed. This optimal configuration was identified only late in the work, and was therefore used to characterize the signal at the output pad only. The rest of the analysis is based on measurements taken with the oscilloscope.

Physiological signal playback With the oscilloscope settings described above, it is possible to characterize the ability of the generator to replay physiological signals down to a few millivolts. The measurements can therefore be performed at the output pad of the signal, characterizing the full analog chain of the generator. The oscilloscope is not precise enough to reach the physiological amplitude range of the EEG, EOG, and EMG signals. However, the ability to reach such amplitudes is characterized by the DC measurements described above. Therefore, if the generator is able to replay signals at amplitudes measurable by the oscilloscope, its ability to generate signals down to physiological amplitudes depends on the DC characterization of the generator. In this way, a complete characterization of the physiological signal generator can be performed. The use of the DC nanovoltmeter to measure the output signal validates the analysis above.

In order to characterize the fidelity of the playback, a protocol is needed to ensure a valid comparison between the generated signal and the original signal. The reference signal chosen is the signal stored in the .bin file on the SD card. In comparison of the original signal from the database, it has been clamped to the 99.5th percentile of the signal to propose a maximal

dynamic range resolution of the useful signal. Hours of signals are recorded and stored on the SD card. However, only short signal samples are sufficient to characterize the generator. For simplicity, it was chosen to record 20-second samples starting from time zero in the file. Initially, an automated benchtest similar to the one used with the DC nanovoltmeter was employed to record data and automate the alignment with the original signal. However, it was observed that substantial noise, on the order of millivolts, was added to the recording when the oscilloscope was connected to the host PC through a USB cable. The recordings were therefore acquired manually on the oscilloscope.

Three operations must be performed before comparing the signals. Firstly, the original signal is resampled to the measurement rate to have the same time basis. Secondly, because the output signal amplitude may not correspond to the original signal amplitude, it must be scaled. Both signals are normalized by dividing them by their rms value after mean removal. Finally, because the recordings are performed manually on the oscilloscope, the recording start time may not correspond to time 0 in the original signal. Time alignment between the two samples is corrected using cross-correlation. Once the samples have the same amplitude and are time-aligned, they can be compared. The metrics used to compare the waveforms are the Pearson correlation coefficient and the ratio between Root Mean Square Error (RMSE) and the RMS of the reference signal. The Pearson correlation coefficient is defined in Equation (3.6), it measures the similarity of the measured signal to the original. The reconstruction error is evaluated using the normalized error, i.e., the RMSE/RMS ratio, the RMSE is defined in Equation (3.7), it computes the average magnitude of the error between the two signals and the RMS is the one of the original signal. The normalized error is used instead of the RMSE alone because the analysis of the output signal is not only done at the original amplitude of the signal. The results are analyzed in Section 3.

$$r = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^N (x_i - \bar{x})^2} \sqrt{\sum_{i=1}^N (y_i - \bar{y})^2}} \quad (3.6)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - y_i)^2} \quad (3.7)$$

where x_i and y_i are the i -th samples of the original and measured signals, respectively, $\bar{x}$ and $\bar{y}$ are the corresponding mean values, and N is the total number of samples.

Validation of the Generator Modes Besides the quantitative characterization of the generator described earlier in this section, the generator includes features that can be qualitatively analyzed. The generator is not only able to play back raw signals from a database, but it can also mix time-aligned artifacts into the cleaned signal, inject artifact samples on demand, and select a specific sleep stage to play, depending on the mode chosen by the user. In order to characterize the ability of the generator to fulfill these roles, one can record signals with mixed artifacts and injected artifacts and compare them qualitatively with the cleaned signal. For the artifact mixing and injection, the procedure is similar to the previous one and consists of recording the first 20 seconds of the signal starting from time 0. Since the output range was characterized in the previous protocols, the validation of the different modes can be performed at the DAC output, as the modes are only firmware-dependent. Specific details about the artifact mixing and injection gain are given in Section 3.3.

2 DC characterization results

This section is dedicated to the analysis of the results obtained with the measurement method and protocols described in section 1.1.1.

2.1 Output range

The output range is determined by the divider settings, i.e., the values of R_1 and R_3 , whose relationship to the code is described in Equation 3.1. The goal of this section is to measure the output range of the generator and compare it to the simulated output range. A curve similar to Figure 2.15 is expected for every channel. Deviations from theoretical values are not problematic as long as the physiological amplitude range of each signal can be achieved (They are defined in Section 1).

First, the measurement protocol described in Section 1.1.1 was applied to the ECG channel. For both R_1 values, the output range obtained is lower than expected for the first codes and then increases to reach a constant factor (≈ 1.25) above the expected value (Figure 3.2). The physiological amplitude range can be reached by the generator for the ECG channel. The measured output range compared to the theoretical output range is shown in Table 3.2 for both R_1 values. The difference between the two can be explained by a wiper resistance that is different from the theoretical value used for predictions and by non idealities in the circuit.




Figure 3.2: Differential output voltage between ECG+ and ECG-, $R_1=200\Omega$, $R_{wiper} = 110\Omega$ for the predicted output, DAC ECG channel at full scale (3.3V at the DAC output).

Table 3.1: Measured vs. theoretical output amplitude range of the ECG differential channel. The minimum output amplitude range corresponds to the maximum value of the digipot and the maximum amplitude output range to the minimum value of the digipot. The theoretical output amplitude range is computed with $R_W = 110\Omega$

R_1	measured min-max [mV]	theoretical min-max [mV]
200 Ω	4.4 - 78.5	3.3 - 165
2 k Ω	0.576 - 6.3	0.4 - 19.3

The measurement protocol was then applied to the single-ended channels for both R_1 values. The results must be further analyzed than the ECG differential channel because of the different topology for these channels. First, when the $1\text{ k}\Omega$ resistance is selected for R_1 , the output range corresponds approximately to the one expected, but a DC offset is observed depending on the channel location, as can be seen from the measured output range in function of the digipot code shown in Figure 3.3. This phenomenon may be a visualization of the voltage drops in the V_{ref} trace discussed in Section 2.2.5.

When the $10\text{ k}\Omega$ resistance is selected for R_1 , the output range corresponds approximately to the one expected. However, depending on the switches configuration, a varying DC offset is observed for each channel. Different switches configuration were tested for the EEG1 channel and the results are shown in Figure 3.4. Removing the offset allowed to compute the measured output range for $R_1 = 10\text{ k}\Omega$ when all switches, besides the one from the measured channel, are open in ADG714 # 2 and closed in ADG714 # 1.

The measured output range compared to the theoretical output range for the EEG1 channel is shown in Table 3.2 for both R_1 values. The minimum output amplitude measured is higher than the theoretical minimum output amplitude but the range of the physiological amplitudes described in Section 1 can still be achieved for each signal type.




Figure 3.3: Differential output voltage between the active output pad and V_{ref} pad with $R_1=1\text{ k}\Omega$, $R_{\text{wiper}} = 110\Omega$ for the predicted output and the DAC output of each channel at full scale. The label correspond to the name of the channel measured.




Figure 3.4: Differential output voltage between the EEG1 output pad and the V_{ref} pad, with $R_1 = 10 \text{ k}\Omega$ and $R_{\text{wiper}} = 110 \Omega$ for the predicted output, and the EEG1 DAC output at full scale. The label indicates the number of switches (other than the one of the measured channel) closed in ADG714 #2 (which selects $R_1 = 10 \text{ k}\Omega$); switches not closed in #2 are closed in #1.

Table 3.2: Measured vs. theoretical output amplitude range of the EEG1 channel. The minimum output amplitude range corresponds to the maximum value of the digipot and the maximum output amplitude range to the minimum value of the digipot. The theoretical output amplitude range is computed with $R_W = 110 \Omega$.

R_1	Measured min-max [mV]	Theoretical min-max [mV]
1 k Ω	0.326 – 14.7	0.390 – 18.9
10 k Ω	0.062 – 1.6	0.041 – 2.0

2.2 Output linearity

Before analyzing the generator output linearity, it is important to ensure that the DAC output itself is linear. The DC voltmeter is used with the positive terminal connected to the channel test point at the RC filter output and the negative terminal connected to ground. Figure 3.5 shows the measured voltage of channel EEG1 as a function of the DAC codes. The output is a linear slope, as expected. Therefore, any non-linearity downstream is caused by the topology of the circuit itself, and not by the DAC. It can be observed that the extreme codes introduce a small error at the output; the signal should therefore be scaled so as to avoid reaching the extreme codes, in order to minimize signal distortion.




Figure 3.5: Upper graph represents the measured voltage of the EEG1 channel at the output of the RC filter in function of the DAC code.

The measurement protocol described in Section 1.1 is applied to measure the generator output linearity. The result is expected to be a downscaled version of Figure 3.5, depending on the divider ratio. The schematic of the divider is shown in Figure 3.1 to visualize the configurations used below.

First, the protocol was applied to measure the ECG differential output. Both R_1 values and a few digipot values were tested. When the $200\ \Omega$ resistance is selected for R_1 , the transfer function is indeed linear and corresponds to the DAC output linearity (Figure 3.6). The transfer function is the same regardless of the digipot value. When the $2\ \text{k}\Omega$ resistance is selected for R_1 , the transfer function depends on the switch configuration. When all switches in ADG #2 are closed, the transfer function is non-linear; however, when all switches in ADG #2 are open and all switches in ADG #1 are closed (apart from the ECG switch), the transfer function is linear (Figure 3.7). As for $R_1 = 200\ \Omega$, the transfer function has a similar shape across the digipot values.

In the best-case configurations, the maximum INL is $40\ \mu\text{V}$ for $R_1 = 2\ \text{k}\Omega$ and $267\ \mu\text{V}$ for $R_1 = 200\ \Omega$, at the extremes of the DAC code range. This error is acceptable for signal integrity, as it occurs at the extremes of the DAC codes and amounts to approximately 6 % of the signal amplitude. The output linearity of the ECG channel is validated for every generator setting except when all switches are closed in the ADG #2.

In the case where the transfer function is non-linear, it is interesting to note that, beyond having a similar shape, the INL remains constant in voltage units across all digipot values, which results in an increasing INL in LSB units as the output range decreases. It can be deduced from this observation that the non-linear transformation does not originate from the digipot, but must appear somewhere else in the circuit. The same protocol was then applied to the channel while the other channels were held at different constant values: all other channels at full DAC scale, and all other channels at zero DAC scale. The transfer function appears to depend on the values of the other channels: when they are at FULL scale, the transfer function is approximately the reflection about the vertical axis of the one obtained when the other channels are at MID scale; when they are at ZERO scale, the

transfer function is more linear (Figure 3.8).




Figure 3.6: Differential output voltage between ECG+ and ECG-, $R_1=200\Omega$, $R_3 \approx 10k\Omega$ (code 0)




Figure 3.7: Comparison of the differential output voltage between ECG+ and ECG- when other ADG#2 switches are closed or open. $R_1=2k\Omega$, $R_3 \approx 10k\Omega$ (code 0).




Figure 3.8: Comparison of differential the output voltage between ECG+ and ECG- for different DAC codes of the other channels. $R_1=2k\Omega$, $R_3 \approx 10k\Omega$ (code 0). All other channel $R_1=10k\Omega$.

The protocol was then applied to measure the single-ended outputs. Both R_1 values and a few digipot values were tested. As with the ECG output, the results differ based on the R_1 value. Moreover, they are related to the results obtained while measuring the output range.

When the $1k\Omega$ resistance is selected for R_1 , the transfer function is linear as long as the other channels are held at constant values other than MID scale. When the other channels are at MID scale, as in the ECG $2k\Omega$ case, the INL remains approximately constant in voltage units across all digipot values, making the transfer function less and less linear as the output range decreases. Figure 3.9 shows the output transfer function of the EEG1 channel as a function of the other channels' values.




Figure 3.9: Comparison of the differential output voltage between EEG1 output pad and V_{ref} pad for different DAC codes of the other channels. $R_1=1k\Omega$, $R_3 \approx 10k\Omega$ (code 0). All other channel $R_1=1k\Omega$.

When the $10\text{ k}\Omega$ resistance is selected for R_1 , the transfer function is linear as long as the other channels are held at constant values other than MID scale **and** as long as at least five of the seven other channels R_1 values are not equal to $10\text{ k}\Omega$. Figure 3.10 shows the output transfer function of the EEG1 channel as a function of the other channels' R_1 values. When these conditions are not met, the transfer function becomes non-linear, introducing distortion in the output signal.

In the best-case configurations, the maximum INL is $6.7\text{ }\mu\text{V}$ for $R_1 = 10\text{ k}\Omega$ and $12\text{ }\mu\text{V}$ for $R_1 = 1\text{ k}\Omega$, both with R_3 . These errors amount to 10.8% and 3.7% of the signal amplitude, respectively. These relative errors are the worst case scenario as they are computed for the lower band of the output amplitude range. They are therefore considered as small enough to not significantly distort the signal in either case. The output linearity of the single-ended channels is therefore validated under the ideal generator settings.




Figure 3.10: Comparison of the differential output voltage between EEG1 output pad and V_{ref} pad for different configurations of the switches. $R_1=1\text{ k}\Omega$, $R_3 \approx 10\text{ k}\Omega$ (code 0). All other channel $R_1=1\text{ k}\Omega$. The label correspond to the number of other switches closed in the ADGG714 #2, that selects $R_1=10\text{ k}\Omega$, if it is not closed in #2, it is closed in #1.

2.3 Noise floor of the generator

The measurement protocol described in Section 1.1 is applied to characterize the noise floor of the generator. The results are summarized in Table 3.3.

First, it was applied to the ECG differential output channel for both values of R_1 and for the minimum and maximum digipot values. When the $200\text{ }\Omega$ resistance is selected for R_1 , the AC RMS noise ranges from approximately $10\text{ }\mu\text{V}$ for the maximum digipot value to approximately $60\text{ }\mu\text{V}$ for the minimum digipot value. When the $2\text{ k}\Omega$ resistance is selected for R_1 , the AC RMS noise is approximately $25\text{ }\mu\text{V}$ regardless of the digipot value.

The protocol was then applied to a single-ended output for both values of R_1 and for the minimum and maximum digipot values. When the $1\text{ k}\Omega$ resistance is selected for R_1 , the AC RMS noise is approximately $8\text{ }\mu\text{V}$ regardless of the digipot value. When the $10\text{ k}\Omega$ resistance is selected for R_1 , the AC RMS noise is approximately $15\text{ }\mu\text{V}$ regardless of the digipot value.

For the ECG channel, these values are acceptable over the whole range; they do not compromise the signal integrity. The lowest measured output amplitude is $576\mu\text{V}$, so the noise accounts for at most 4.34% of the signal amplitude in the worst case. For the single-ended channels, however, the lowest measured output amplitude is $62\mu\text{V}$, so the noise can account for up to 24.2% of the signal amplitude at the lowest output amplitudes, which compromises signal integrity at the extreme low end of the range.

Table 3.3: AC RMS noise measurements for ECG and single-ended outputs.

Output	R_1 [Ω]	Digipot value	AC RMS noise [μV]	Minimum measured output amplitude [μV]
ECG	200	Maximum	~ 10	78500
		Minimum	~ 60	4400
	2 k	Min / Max	~ 25	576
Single-ended	1 k	Min / Max	~ 8	326
	10 k	Min / Max	~ 15	62

2.4 Discussion about the DC characteristics of the generator

The numerous results analyzed above tell us about the ability of the generator mentioned at the beginning of this section.

First, the generator's DC characteristics depend on its settings. The ideal settings, i.e., those under which the generator's characteristics are at their best, are characterized by two conditions for the single ended channels (only the second condition applies to the ECG channel). The first one is that the channels that are not being measured are held at a fixed DAC output voltage different from MID scale (i.e., away from V_{ref}), and the second one is that the switches of the second ADG714 are not all closed simultaneously. Under these settings, the generator verify several key aspects: the ability to reach the targeted output amplitudes, the transfer of the signal without distortion, and an acceptable worst-case SNR. The ideal characteristics of the generator are presented in Table 3.4.

Table 3.4: Ideal DC characteristics of the physiological signal generator. The worst-case SNR is computed as $20 \log_{10}(V_{\text{out,min}}/V_{\text{noise,max}})$. The MAX INL/ampl ratio is computed at the minimum output amplitude.

Characteristics	ECG (differential)		Single-ended	
	$R_1 = 200\Omega$	$R_1 = 2\text{ k}\Omega$	$R_1 = 1\text{ k}\Omega$	$R_1 = 10\text{ k}\Omega$
Output range [mV]	4.4 – 78.5	0.576 – 6.3	0.326 – 14.7	0.062 – 1.6
Max INL [μV]	267	40	12	6.7
MAX INL / min ampl [%]	6.1	6.9	3.7	10.8
AC RMS noise [μV]	10 – 60	~ 25	~ 8	~ 15
Worst-case SNR [dB]	37.3	27.3	32.2	12.3

Second, the difference between the results of the ECG channel and the single-ended channels is explained by their respective topologies. The ECG output is a true differential output, independent of V_{ref} . The single-ended outputs, in contrast, swing around the local V_{ref} at the bottom of their divider, while the DC voltmeter references the signal to the V_{ref} at the output pad. Because these two points on the V_{ref} node are not at the same potential

when inter-channel currents are present, the single-ended outputs are sensitive to coupling that the ECG output rejects.

Moreover, the results obtained under non-ideal settings reveal a design issue explored in Section 2.2.5. Due to the analog reference topology, the different channels are not independent, they are coupled through the V_{ref} node, and cross-channel coupling currents modulate the output voltage of each channel. As shown above, the non-idealities of a given channel depend on the settings of the other channels and on its own injecting current. As an additional test to confirm these observations, probing the channel output with the reference probe placed at the local V_{ref} node (at the bottom of R4) yields non-idealities of the same nature as probing against the V_{ref} pad, though with a slightly different shape (Figure 3.11). This excludes a static DC offset along the V_{ref} trace between the divider and the pad as the only cause of the non-idealities. The coupling occurs at the V_{ref} node itself, consistent with currents from the channel itself and from neighboring channels modulating the node through its non-zero impedance. The two probe positions capture the same underlying coupling but measure the modulated node at different points. This inter-channel coupling introduces non-idealities in the output signal. The magnitude of this effect is difficult to estimate under realistic playback configurations. However, based on the analysis above, it can be minimized by configuring the generator appropriately.

Finally, the results show that when the ADG #2 switch is closed for all channels, the output deviates significantly from the ideal. Closing all these switches simultaneously should therefore be avoided in order to minimize signal degradation. This does not result from the same coupling phenomenon described above, as it affects not only the single-ended channels but also the differential channel. Moreover, beyond introducing distortion, it prevents the output from reaching the range set by the divider ratio. It was therefore concluded that this phenomenon originates from the ADG #2 itself when most of its switches are closed, possibly due to a device-level issue introducing leakage between channels. This remains a hypothesis. Confirming it would require either a closer characterization of the ADG #2 behavior in this specific configuration or replacing the device to check whether the effect persists. This was not done in the present work, as the effect does not prevent the generator from operating and the available time was limited.

The effect of these behaviors on the ability of the generator to faithfully reproduce physiological signals is assessed in the next Section.

Regarding a possible voltage difference between V_{ref} at the node and V_{ref} at the pad, this would manifest as a DC offset on the measured signal, since the difference is not seen as a common-mode voltage by the acquisition system. This is not a concern in practice, as most recording devices include a high-pass filter to remove the DC component of the signal. Therefore, even if the generator introduces a DC offset, the signal itself is not compromised.




Figure 3.11: Comparison of the differential output voltage between EEG1 and Vref at the pad and at the bottom of R_4 . $R_1=1k\Omega$, $R_1 \approx 10k\Omega$ (code 0), the other channels settings are : DAC at MID scale, $R_1 = R_{1a}$ and digipot $\approx 110\Omega$ (code 255).

3 Dynamic characterization results

This section is dedicated to the analysis of the results obtained with the measurement methods and protocols described in Section 1.1.2. Before proceeding with the analysis, it is important to determine the supply noise floor level in order to identify the output amplitude range down to which relevant measurements can be performed using the protocol with the oscilloscope described in Section 1.1.2.

3.1 Supply noise

The oscilloscope is used to determine the supply noise. In high-resolution and AC mode, the oscilloscope probes the V_{ref} pad and one output of the DAC, coded at 0, in order to observe the noise present at each node. The first measurement provides information about the noise present at the V_{ref} node, and the second about the noise present at the DAC output. The latter is attenuated by the divider downstream, whereas the former determines the noise present at the output of every channel, as it couples through the divider.

The comparison of the two supplies was clearly in favor of the Margherita wall supply, providing between 200 and 300 μV of noise. The RC filter before the V_{ref} buffer is intended to filter the supply noise with a -3 dB cut-off at 3.2 Hz; however, no noise filtration is observed. This might be due to the fact that the 10 μF capacitor used is a 25 V aluminum electrolytic capacitor, which is not suited for a 1.65 V application. A more adequate ceramic capacitor would have allowed to achieve a lower supply noise in the analog reference.

This noise level is too high to properly measure signals down to the physiological amplitude range in a single-ended oscilloscope measurement setup; however, the measurements performed at a higher scale are representative of what happens at a lower scale, thanks to the DC characterization. Moreover, using the adequate measurement setup, i.e. a differential measurement setup, the supply noise should be rejected as a common-mode voltage between the V_{ref} pad and the active pad.

Supply	Measurement point	Noise RMS [mV]
TENMA	DAC output	1.35
	Vref pad	1.35
Margherita supply	DAC output	0.31
	Vref pad	0.30

Table 3.5: Power supply noise measurements at DAC output and Vref pad.

3.2 Physiological signal playback

Despite the DC characteristics of the generator not being ideal for every setting, its ability to generate signals from digital data can still be characterized.

To assess the generator’s ability to faithfully replay a physiological signal, the Pearson correlation is used to quantify the similarity between the original and the measured signal. The higher the Pearson coefficient, the more similar the two signals; a coefficient of 1 corresponds to a perfect linear relationship between the two signals. Preserving the general morphology and events of the signal is already a relevant target to achieve. The ideal is a coefficient of 1, the aim is to get as close as possible to this number, confirming that the signal’s morphology and events are reproduced. The RMSE/RMS ratio, i.e, the normalized error, quantifies the relative error between the generated and original signals. A lower ratio indicates a smaller error.

Signal output characterization using the oscilloscope: The protocol described in Section 1.1.2 was applied to two signals, EEG1 and ECG, representing each of the two analog chains. It was applied at every stage of the chain: the DAC output, the intermediate divider stage, and the output pad. The ability to reach the target physiological amplitude while maintaining a high correlation is not characterized by this protocol, since the setup cannot cancel the supply noise. Therefore, the major noise contribution in this measurement setup is not the generator noise or channel coupling phenomenon but the supply noise.

It was first applied to the ECG channel. The correlation coefficient is almost perfect, with $r = 0.99$ at the DAC output and around $r = 0.85$ at the output pad for divider settings of $R_1 = 200\Omega$ and $R_3 = 10k\Omega$, giving an output amplitude of approximately 5 mV (Figure 3.12). The decrease in the correlation coefficient is due to the noise becoming proportionally more significant.




Figure 3.12: Comparison of the original signal with the output voltage measured at the ECG+ pad, for $R_1=200\Omega$, $R_1 = 10k\Omega$ (code=0). The original signal is the ECG channel.

It was then applied to the EEG1 channel. The correlation coefficient is almost perfect, with $r = 0.99$ at the DAC output. It is constant, around $r = 0.66$, at the intermediate and output stages for divider settings of $R_1 = 1\text{ k}\Omega$ and $R_3 \approx 110\text{ }\Omega$, giving an output amplitude of approximately 10 mV (Figure 3.13). The correlation coefficient decreases because the noise becomes proportionally more important.




Figure 3.13: Comparison of the original signal with the output voltage measured at the EEG1 pad, for $R_1=1\text{ k}\Omega$, $R_3= R_W \approx 110\text{ }\Omega$ (code=255). The original signal is the raw EEG1 channel.

Signal output characterization using the DC nanovoltmeter: The protocol using the DC nanovoltmeter described in Section 1.1.2 was applied to the ECG and EEG1 channels, as they are each representative of one of the two topologies. The alternative measurement setup proved to be more than just an additional verification of the differential configuration. In practice, it can sample the signal at a sufficient rate to faithfully reproduce the waveform, allowing a more thorough validation of the generator.

In the ideal configuration of the generator for single ended channels, i.e., one active channel while the others are held at constant values different from V_{ref} and not all switch from ADG714 # 2 closed simultaneously, the minimum output amplitude setting was first tested on the EEG1 channel, and the results are consistent with the DC characteristics of the generator. The signal appears noisy and loses fidelity compared with the original, as expected from the $15\text{ }\mu\text{V}$ noise floor of the generator at this setting. However, the signal remains distinguishable with a correlation coefficient of 0.7, as shown in Figure 3.14. Figure 3.15 shows the output signal at an amplitude about three times that of the original signal, with a correlation coefficient close to 1. The protocol was then applied to the ECG channel under its ideal configuration, i.e., not all the switches from ADG # 2 are closed simultaneously. The generated signal at physiological amplitude has a correlation coefficient close to 1, as can be seen in Figure 3.16. This validates the ability of the generator to faithfully replay signals at the intended output amplitude under the ideal generator settings.




Figure 3.14: Upper graph) Measured output voltage between EEG1 pad and V_{ref} pad. $R_1=10\text{k}\Omega$, $R_3\approx 10\text{k}\Omega$ (code=0). Lower graph) Comparison of the measured EEG1 signal with the original signal.




Figure 3.15: Upper graph) Measured output voltage between EEG1 pad and V_{ref} pad. $R_1=10\text{k}\Omega$, $R_3\approx 1\text{k}\Omega$ (code=220). Lower graph) Comparison of the measured EEG1 signal with the original signal.




Figure 3.16: Upper graph) Measured differential output voltage between ECG+ and ECG-. $R_1=2k\Omega$, $R_3\approx 1k\Omega$ (code=220). Lower graph) Comparison of the measured ECG signal with the original signal.

The output signal of the generator was also measured under the normal operating settings, i.e., when all channels are actively replaying signals. The results show, as expected from the DC characterization, that inter-channel coupling becomes a concern for signal fidelity below the millivolt range, where the signal becomes too noisy. Above an amplitude of one millivolt, the signal output remains faithful to the original even under the normal settings. A comparison of the EEG1 output signal between the ideal and normal settings is shown in Figure 3.17, for an ideal output amplitude of approximately $400\mu V$.




Figure 3.17: Comparison of the measured EEG1 output for two configuration of the generator. Configuration 1: other channels are playing; configuration 2: other channels remain at a constant value. $R_1=1k$, $R_3\approx 10k\Omega$ (code=0).

3.3 Validation of the Generator Modes

The goal of this section is to validate the different modes of the generator, i.e., the artifact mixing mode, the artifact injection mode, and its ability to select specific sleep stages. The procedure described in Section 1.1.2 is applied in order to validate the different modes.

3.3.1 Artifacts mixing

The mixing of artifacts into the cleaned signal is properly handled by the generator (Figure 3.18). The ECG artifact is identified by the negative spikes at regular intervals. The EOG artifact is more easily distinguished when large blinks occur, which is not the case in this sample; it can still be detected here through the slight baseline difference between the two signals, which corresponds to the EOG contribution to the EEG channel. The effect of the gain setting is also validated; all gains are equal to one in the first part of the recording, then the ECG gain is increased to three, and the negative spikes increase accordingly.




Figure 3.18: Artifact mixing in the cleaned signal with $\alpha_{ecg} = 1$ from time 0 to 13 seconds and then $\alpha_{ecg} = 3$ from 13 seconds until the end. The original signal is the clean EEG1 channel.

3.3.2 Artifacts injection

The injection of artifact samples into the signal is properly handled by the generator. Figure 3.19 shows the injection of an ECG burst with a gain of 3 into the cleaned signal starting at 12 seconds. The ECG artifact is clearly distinguishable with negative spikes at constant intervals that correspond to the R-peaks of the ECG, as seen by the EEG1 channel. Figure 3.20 shows the injection of an EOG blink with a gain of 0.5 into the cleaned signal at 15 seconds.




Figure 3.19: Injection of an ECG burst with a gain of 3 into the cleaned EEG1 signal. The signal is measured at the DAC output. The original signal is the clean EEG1 channel. Measured at the RC filter output.




Figure 3.20: Injection of an EOG blink artifact with a gain of 0.5 into the cleaned EEG1 signal. The signal is measured at the DAC output. The original signal is the clean EEG1 channel. Measured at the RC filter output.

3.3.3 Sleep stage playback

The figures below represent the EEG1 signal at the output of the RC filter after the DAC for different sleep stages. As for the artifact mixing and injection protocol, the signal was measured at the RC filter output instead of being measured at the generator output, since the goal here is to verify that the generator can select a specific sleep stage, not to characterize the output (which has been done above). The measured signal is compared to the EEG1 raw signal stored on the SD card. The principle used to compare the original and the measured signal is the same as above, with one difference: the recorded signal does not start at time zero but at the beginning of the selected stage, located by the firmware. The Figures 3.21 to 3.25 show that the sleep stage playback is validated.




Figure 3.21: Playback of one wake stage. Measured at the RC filter output.




Figure 3.22: Playback of one N1 stage. Measured at the RC filter output.




Figure 3.23: Playback of one N2 stage. Measured at the RC filter output.




Figure 3.24: Playback of one N3 stage. Measured at the RC filter output.




Figure 3.25: Playback of one REM stage. Measured at the RC filter output.

3.4 Discussion about the dynamic characteristics of the generator

The analysis of the results in Section 3.2 validates the ability of the generator to faithfully replay signals at its outputs and confirms the results and analysis presented in Section 2.4. With only one single-ended channel active, signal fidelity is excellent down to output amplitudes of a few hundred microvolts and remains acceptable at the minimum scale. When all channels are used simultaneously, the inter-channel coupling degrades the output signal quality and becomes concerning below the millivolt range for single-ended channels. The ECG channel, however, is able to replay signals without being affected by inter-channel coupling at any range, and is therefore validated for all generator settings. The results of Section 3.3 further validate the generator's ability to play back signals in its three modes (raw playback, artifact mixing, and artifact injection), with the additional option of selecting a specific sleep

stage. Moreover, the intrinsic noise of the generator could be considered an artifact in itself, present in the output signal regardless of the mode. Depending on the output amplitude, the noise contribution to the signal could be computed and, together with the contributions from the mixed or injected artifacts, the total artifact level could be communicated to the user. This feature is not implemented in the current design, but it could easily be added in a future version of the firmware.

The DC offset appearing at the output on single ended channels is not a concern, as it does not affect the signal integrity and can be removed thanks to a high pass filter at the acquisition if needed.

An important result of the dynamic characterization of the generator is that the supply noise is fully rejected by the differential topology of the output, for both the true differential channel (ECG) and the single-ended channels referenced to V_{ref} . In the DC characterization, the low-pass filter and the long integration time were able to filter out the mains interference from the supply. In the dynamic protocol, neither is active, so the measured signal is the one actually generated, without any noise removal beyond what is provided by the topology itself. This is a key feature of the generator; the only noise concern is the generator's intrinsic noise, and the choice of supply is therefore not a significant constraint for the user.

Chapter 4

Conclusion and perspectives

1 Conclusion

The objectives of this master thesis were to design and implement an accessible physiological signal generator that replays real signals from a database, injects configurable artifacts, and supports dry-electrode connections. The work began with an introduction to the physiological signals relevant to the PSG, followed by a review of the state of the art in head phantoms and biosignal generators. The design of the generator was then divided into three main parts. Some design choices in one part influenced the others, while certain blocks are independent and can be replaced or improved without affecting the rest, giving the generator a flexible, modular design. The signal processing part was designed first, upstream of the other parts. The hardware design then began, and once the PCB was functional, the firmware was implemented. Once everything was working as intended, the generator was characterized. The analysis of the DC characteristics and the output signal both established the generator's specifications and revealed its limitations. The generator designed in this master thesis costs approximately 200 euros.

The objectives of this work were met. The generator is accessible, can replay signals from human recordings, can add artifacts in a controllable manner, and can be connected to dry electrodes. It can faithfully replay physiological signals without limiting constraints for the ECG channel. For the single-ended channels (EEG, EOG, EMG), however, two configurations allow faithful signal replay: either a single channel at physiological amplitudes, or all channels simultaneously at amplitudes higher than the physiological range. Moreover, the use of a wall supply, not ideal in theory for a low-noise circuit, was found not to affect the output signal, thanks to the generator's design, which causes the supply noise to appear as a common-mode signal at the acquisition system, where it is rejected. This is a particularly useful property, as it makes an ultra-low-voltage generator **largely insensitive to supply noise** and compatible with any reasonable power supply, without concerns about battery lifetime during playback. The dominant limitation lies elsewhere. The generator's ability to replay signals at physiological amplitudes across all eight channels **simultaneously** is constrained by **inter-channel coupling**, which introduces excessive distortion at sub-millivolt amplitudes. This does not impact the ECG differential channel.

The choice of using an SD card to store the signal data makes the generator an independent device, controlled only via a computer, and gives the user the option to add new data. It also offers the possibility to derive different generator modes based on the same

architecture.

The main design issue responsible for inter-channel coupling is the shared analog reference node across all channels with a non-zero impedance in addition with small value resistances connected to the analog reference. One way to mitigate this effect is to replace the trace with a copper pour to reduce the impedance and to use higher resistances values connected to the analog reference to reduce the current sources.

2 Perspectives

Several directions can be explored to extend and improve this work. These directions are separated into hardware improvements, software improvements, and testing and future development.

Hardware improvements: First, implementing the suggested fix to the V_{ref} topology would enable simultaneous replay across all channels and allow the multimodal capability of the generator to be fully validated, whatever the output amplitude. Other hardware improvements could be made to consolidate and refine the design, including integrating all components onto a single board.

Software improvements: The signal-processing pipeline used for the artifact mixing and injection feature is a preliminary version of what could be further developed. A deeper investigation of this topic could enable the injection of more relevant artifacts tailored to the validation requirements or to test specific failure modes (automatic artifact rejection algorithms for instance) of acquisition systems. Artifact injection could also be extended to non-physiological artifacts, such as electrode motion artifacts, which are not currently included in a controllable manner.

Testing and future development: Overall, the work demonstrates that accessible, multimodal, hardware-level PSG validation with configurable artifact injection is feasible and provides a documented basis for future versions. The generator presented in this work is functional for testing contactless electrodes. All that remains is to put it to use. Future work could explore new generator designs that improve their resemblance to the human body, such as a head-shaped generator or one that simulates human skin impedance at the output pad.

Appendices

Appendix A Generator schematics

This appendix presents the KiCad schematics of the generator board. An overview of the complete KiCad schematic file is shown in Figure 4.1. Figures 4.2 to 4.10 show the schematics of the different parts that constitute the generator circuit. The schematics offer an assembled visualization of the circuit parts described in Section 2.1. The schematics present the fixes for the different floating pins described in Section 2.2.7. One should note that the labels used for the analog signal from the RC filter output to the generator output are global label.

When designing the generator circuit, it was chosen to implement LEDs on the board to visualize when a channel was active. However, due to time constraints and the optional character of this implementation, the firmware does not include this feature on the current version of the generator.




Figure 4.1: Overview of the generator schematic. The LEDs were not




Figure 4.2: Connector to the Margherita board. The connectors cables are providing the power supply and the I2C and SPI communication protocols from the Margherita board.




Figure 4.3: The 8 channels DAC7578.




Figure 4.6: Schematic of the octal switches. Each channel has a path in each switch; the two paths are in parallel and are used to select the first resistor of the divider (R_1). For each channel, the signal entering the switch is labeled with *channel_filt*, originating from the buffered output of the RC filter. The signal leaving the switch is labeled with *SR_{value}_channel*.




Figure 4.7: EEG1 and EEG3 output stages (single ended channels). For each channel, the input of the divider corresponds to the signal that went either through $R_1=1k\Omega$ or $R_1=10k\Omega$. The signal is then attenuated through the two stage divider and the output is buffered. The SDI input of the digipot is the output of the previous device on the SPI chain. The RC in the feedback loop of the output buffer has no effect (Section 2.1.3).




Figure 4.8: EEG2 and EEG4 output stages (single ended channels). For each channel, the input of the divider corresponds to the signal that went either through $R_1=1\text{k}\Omega$ or $R_1=10\text{k}\Omega$. The signal is then attenuated through the two stage divider and the output is buffered. The SDI input of the digipot is the output of the previous device on the SPI chain. The RC in the feedback loop of the output buffer has no effect (Section 2.1.3).




Figure 4.9: EOG1 and EOG2 output stages (single ended channels). For each channel, the input of the divider corresponds to the signal that went either through $R_1=1\text{k}\Omega$ or $R_1=10\text{k}\Omega$. The signal is then attenuated through the two stage divider and the output is buffered. The SDI input of the digipot is the output of the previous device on the SPI chain. The RC in the feedback loop of the output buffer has no effect (Section 2.1.3).






Figure 4.10: EMG (single ended channel) and ECG (differential channel) output stages. For the EMG channel, the input of the divider corresponds to the signal that went either through $R_1=1\text{k}\Omega$ or $R_1=10\text{k}\Omega$ or through $R_1=200\Omega$ or $R_1=2\text{k}\Omega$ for the ECG channel. The signal is then attenuated through the two stage divider and the output is buffered. The SDI input of the digipot is the output of the previous device on the SPI chain. The RC in the feedback loop of the output buffer has no effect (Section 2.1.3). The RC in the feedback loop of the inverter of the ECG channel is effective.

Appendix B Generator PCB layout


This appendix presents the complete PCB layout of the generator board. The board is a four layer design, each layer is shown in the Figure below.




(a) Front copper layer



(b) First inner copper layer



(c) Second inner copper layer



(d) Back copper layer

Figure 4.11: PCB layout of the generator board, showing each of the four copper layers individually.

Appendix C Margerhita board

This appendix presents the Microsys module that supports the nucleo h745zi-q board, the SD card interface and the 3.3 V rail. The nucleo board is connected at the back of the board.




Figure 4.12: Margherita board embedding the nucleo h745zi-q, the SD card module and the 3.3V rail. 1) SD card module. 2) Power supply connector (GND = Yellow, 3.3V=green). 3) I2C and 3.3V connector to the generator. 4) SPI connector to the generator.

Appendix D User interface

This appendix describes the commands that the user can use to interact with the generator. Providing a user interface makes the generator more user-friendly, allowing the user to control the playback mode, set the mixing coefficients, and choose which artifact to inject without reflashing the firmware. The commands are presented in Table 4.1.

Table 4.1: User commands exposed over the UART interface. $\langle X \rangle$: {W, 1, 2, 3, R}; $\langle \text{art} \rangle$: {ecg, eog}; channels $\langle \text{ch} \rangle$: [0,7]; $\langle \text{code} \rangle$: [0,255]; $\langle \alpha \rangle$: floating-point; $\langle \text{byte} \rangle$: 8-bit hexadecimal value.

Group	Command	Effect
File System	ls	list file content of the SD card
	lib open $\langle \text{lib.bin} \rangle$	open artifact library and read its index
	lib list	print artifacts samples in library
	lib close	close library
Playback	play raw $\langle \text{raw.bin} \rangle$	start mode 3.3.1
	play art $\langle \text{clean.bin} \rangle$ $\langle \text{ecg.bin} \rangle$ $\langle \text{eog.bin} \rangle$	start mode 3.3.2
	play lib $\langle \text{clean.bin} \rangle$ $\langle \text{lib.bin} \rangle$	start mode 3.3.3 (opens lib.bin)
	play ... stage $\langle X \rangle$ $\langle \text{Score.txt} \rangle$	start mode 3.3.4, ...requires arguments from one of the three modes
	stop	stop current playback
Artifact	set $\langle \text{art} \rangle$ $\langle \text{ch} \rangle$ $\langle \alpha \rangle$	set α_c^{art} for one channel
	setall $\langle \text{art} \rangle$ $\langle \alpha \rangle$	set the coefficient for all channels
	inject $\langle \text{id} \rangle$ [loop] [gain $\langle g \rangle$]	inject artifact sample $\langle \text{id} \rangle$ from the library
Div ratio	chain digiall $\langle \text{code} \rangle$	set all digipots to the corresponding $\langle \text{code} \rangle$ value
	chain digipot $\langle \text{ch} \rangle$ $\langle \text{code} \rangle$	set digipot $\langle \text{ch} \rangle$ to the corresponding $\langle \text{code} \rangle$ value
	chain mux $\langle \text{byte1} \rangle$ $\langle \text{byte2} \rangle$	set ADG714s raw bytes
Help	help	print the command summary
Reset	reset	reset the stm32

Bibliography

- [1] Ohayon, Maurice M. “Epidemiological Overview of Sleep Disorders in the General Population”. In: *Sleep Medicine Research* 2.1 (2011), pp. 1–9.
- [2] N. AlGhanim, V. R. Comondore, J. Fleetham, C. A. Marra, and N. T. Ayas. “The Economic Impact of Obstructive Sleep Apnea”. In: *Lung* 186.1 (Feb. 2008), pp. 7–12.
- [3] S. R. Pandi-Perumal, D. W. Spence, and A. S. BaHammam. “Polysomnography: An Overview”. In: *Primary Care Sleep Medicine: A Practical Guide*. Ed. by J. F. Pagel and S. R. Pandi-Perumal. New York, NY: Springer New York, 2014, pp. 29–42.
- [4] J. V. Rundo and R. Downey. “Chapter 25 - Polysomnography”. In: *Handbook of Clinical Neurology*. Ed. by K. H. Levin and P. Chauvel. Vol. 160. Elsevier, Jan. 2019, pp. 381–392.
- [5] C. Iber, S. Ancoli-Israel, A. Chesson, and S. F. Quan. *The AASM Manual for the Scoring of Sleep and Associated Events: Rules, Terminology and Technical Specifications*. 1st Edition. Westchester: American Academy of Sleep medicine, 2007.
- [6] R. K. Malhotra and A. Y. Avidan. “Sleep Stages and Scoring Technique”. In: *Atlas of Sleep Medicine*. Elsevier, 2014, pp. 77–99.
- [7] D. J. Creel. “Chapter 33 - The Electrooculogram”. In: *Handbook of Clinical Neurology*. Ed. by K. H. Levin and P. Chauvel. Vol. 160. Elsevier, Jan. 2019, pp. 495–499.
- [8] N. Butkov and S. A. Keenan. “An Overview of Polysomnographic Technique”. In: *Sleep Disorders Medicine*. Ed. by S. Chokroverty. New York, NY: Springer New York, 2017, pp. 267–294.
- [9] J. Francis. “ECG Monitoring Leads and Special Leads”. In: *Indian Pacing and Electrophysiology Journal* 16.3 (May 2016), pp. 92–95.
- [10] B. Yi, H. Cao, X. Song, S. Zhao, W. Guo, and M. Li. “How to Identify the Take-over Criticality in Conditionally Automated Driving? An Examination Using Drivers’ Physiological Parameters and Situational Factors”. In: *Transportation Research Part F: Traffic Psychology and Behaviour* 85 (Feb. 2022), pp. 161–178.
- [11] M. Hammad, A. Maher, K. Wang, F. Jiang, and M. Amrani. “Detection of Abnormal Heart Conditions Based on Characteristics of ECG Signals”. In: *Measurement* 125 (Sept. 2018), pp. 634–644.
- [12] H. Tam and J. G. Webster. “Minimizing Electrode Motion Artifact by Skin Abrasion”. In: *IEEE Transactions on Biomedical Engineering* BME-24.2 (Mar. 1977), pp. 134–139.
- [13] J. A. Urigüen and B. Garcia-Zapirain. “EEG Artifact Removal—State-of-the-Art and Guidelines”. In: *Journal of Neural Engineering* 12.3 (June 2015), p. 031001.
- [14] X. Jiang, G.-B. Bian, and Z. Tian. “Removal of Artifacts from EEG Signals: A Review”. In: *Sensors* 19.5 (Feb. 2019), p. 987.
- [15] J. C. Huhta and J. G. Webster. “60-Hz Interference in Electrocardiography”. In: *IEEE Transactions on Biomedical Engineering* BME-20.2 (Mar. 1973), pp. 91–101.

- [16] C. Levkov, G. Mihov, R. Ivanov, I. Daskalov, I. Christov, and I. Dotsinsky. “Removal of Power-Line Interference from the ECG: A Review of the Subtraction Procedure”. In: *BioMedical Engineering OnLine* 4.1 (Dec. 2005), p. 50.
- [17] K. T. Sweeney, T. E. Ward, and S. F. McLoone. “Artifact Removal in Physiological Signals—Practices and Possibilities”. In: *IEEE Transactions on Information Technology in Biomedicine* 16.3 (May 2012), pp. 488–500.
- [18] E.-R. Symeonidou, A. Nordin, W. Hairston, and D. Ferris. “Effects of Cable Sway, Electrode Surface Area, and Electrode Mass on Electroencephalography Signal Quality during Motion”. In: *Sensors* 18.4 (Apr. 2018), p. 1073.
- [19] D. F. Stegeman, D. Dumitru, J. C. King, and K. Roeleveld. “Near- and Far-Fields: Source Characteristics and the Conducting Medium in Neurophysiology:” in: *Journal of Clinical Neurophysiology* 14.5 (Sept. 1997), pp. 429–442.
- [20] J. Malmivuo and R. Plonsey. *Bioelectromagnetism Principles and Applications of Bioelectric and Biomagnetic Fields*. Oxford University Press, Oct. 1995.
- [21] S. Arnau, F. Sharifian, E. Wascher, and M. F. Larra. “Removing the Cardiac Field Artifact from the EEG Using Neural Network Regression”. In: *Psychophysiology* 60.10 (Oct. 2023), e14323.
- [22] G. Dirlich, L. Vogl, M. Plaschke, and F. Strian. “Cardiac Field Effects on the EEG”. In: *Electroencephalography and Clinical Neurophysiology* 102.4 (Apr. 1997), pp. 307–315.
- [23] R. Croft and R. Barry. “Removal of Ocular Artifact from the EEG: A Review”. In: *Neurophysiologie Clinique/Clinical Neurophysiology* 30.1 (Feb. 2000), pp. 5–19.
- [24] M. Iwasaki, C. Kellinghaus, A. V. Alexopoulos, R. C. Burgess, A. N. Kumar, Y. H. Han, H. O. Lüders, and R. J. Leigh. “Effects of Eyelid Closure, Blinks, and Eye Movements on the Electroencephalogram”. In: *Clinical Neurophysiology* 116.4 (Apr. 2005), pp. 878–885.
- [25] M. Fatourechi, A. Bashashati, R. K. Ward, and G. E. Birch. “EMG and EOG Artifacts in Brain Computer Interface Systems: A Survey”. In: *Clinical Neurophysiology* 118.3 (Mar. 2007), pp. 480–494.
- [26] D. F. Stegeman, J. H. Blok, H. J. Hermens, and K. Roeleveld. “Surface EMG Models: Properties and Applications”. In: *Journal of Electromyography and Kinesiology* 10.5 (Oct. 2000), pp. 313–326.
- [27] W. O. Tatum. *Atlas of Artifacts in Clinical Neurophysiology*. New York: Demos, 2019.
- [28] I. Goncharova, D. McFarland, T. Vaughan, and J. Wolpaw. “EMG Contamination of EEG: Spectral and Topographical Characteristics”. In: *Clinical Neurophysiology* 114.9 (Sept. 2003), pp. 1580–1593.
- [29] T. J. Collier, D. B. Kynor, J. Bieszczad, W. E. Audette, E. J. Kobylarz, and S. G. Diamond. “Creation of a Human Head Phantom for Testing of Electroencephalography Equipment and Techniques”. In: *IEEE Transactions on Biomedical Engineering* 59.9 (Sept. 2012), pp. 2628–2634.
- [30] H. McCann, G. Pisano, and L. Beltrachini. “Variation in Reported Human Head Tissue Electrical Conductivity Values”. In: *Brain Topography* 32.5 (Sept. 2019), pp. 825–858.
- [31] P. Akor, G. Enemali, U. Muhammad, J. Crowley, M. Desmulliez, and H. Larijani. “A Cost-Effective 3D-Printed Conductive Phantom for EEG Sensing System Validation: Development, Performance Evaluation, and Comparison with State-of-the-Art Technologies”. In: *Sensors* 25.16 (Aug. 2025), p. 4974.
- [32] W. E. Audette, J. Bieszczad, L. V. Allen, S. G. Diamond, and D. B. Kynor. *Design and Demonstration of a Head Phantom for Testing of Electroencephalography (EEG) Equipment*. Technical Report. Hanover, NH, USA: Creare, Inc., 2020.

- [33] A. Y. Owda and A. J. Casson. “Investigating Gelatine Based Head Phantoms for Electroencephalography Compared to Electrical and Ex Vivo Porcine Skin Models”. In: *IEEE Access* 9 (2021), pp. 96722–96738.
- [34] G. B. Tseghai, B. Malengier, K. A. Fante, and L. Van Langenhove. “A Long-Lasting Textile-Based Anatomically Realistic Head Phantom for Validation of EEG Electrodes”. In: *Sensors* 21.14 (July 2021), p. 4658.
- [35] S. M. Peterson and D. P. Ferris. “Combined Head Phantom and Neural Mass Model Validation of Effective Connectivity Measures”. In: *Journal of Neural Engineering* 16.2 (Apr. 2019), p. 026010.
- [36] A. S. Oliveira, B. R. Schlink, W. D. Hairston, P. König, and D. P. Ferris. “Induction and Separation of Motion Artifacts in EEG Data Using a Mobile Phantom Head Device”. In: *Journal of Neural Engineering* 13.3 (May 2016), p. 036014.
- [37] M. P. Powell, J. Anso, R. Gilron, N. R. Provenza, A. B. Allawala, D. D. Sliva, K. R. Bijanki, D. Oswalt, J. Adkinson, N. Pouratian, S. A. Sheth, W. K. Goodman, S. R. Jones, P. A. Starr, and D. A. Borton. “NeuroDAC: An Open-Source Arbitrary Biosignal Waveform Generator”. In: *Journal of Neural Engineering* 18.1 (Feb. 2021), p. 016010.
- [38] D. Haci, Y. Liu, and T. G. Constandinou. “32-Channel Ultra-Low-Noise Arbitrary Signal Generation Platform for Biopotential Emulation”. In: *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*. Baltimore, MD, USA: IEEE, May 2017, pp. 1–4.
- [39] Fluke Biomedical. *The ProSim 8 and ProSim 8P Vital Signs Patient Simulator*. <https://www.flukebiomedical.com/products/patient-simulators/prosim-8-and-prosim-8p-vital-signs-patient-simulator> (visited on 05/01/2026).
- [40] WhaleTeq. *SECG 5.0 AIOMulti Vital Sign Simulator*. <https://www.whaleteq.com/en/product/1/9-ecg-ekg-testing/view1-secg-50-aio-multi-vital-sign-simulator> (visited on 02/26/2026).
- [41] Datrend Systems Inc. *vPad-A1*. <https://www.datrend.com/product/vpad-a1/> (visited on 05/01/2026). Aug. 2017.
- [42] WhaleTeq. *SEEG 100 | EEG Performance Tester*. <https://www.whaleteq.com/en/product/2/14-eeeg-testing/view15-seeg-100> (visited on 02/26/2026).
- [43] Clara Guenter. *Introducing LabSim : A Signal Generator for Troubleshooting*. <https://pressrelease.brainproducts.com/lablim/> (visited on 05/01/2026). Dec. 2019.
- [44] E. Karakulak. “ARM MCU-Based Experimental EEG Signal Generator Using Internal DAC and PWM Outputs”. In: *Gazi University Journal of Science* 35.3 (Sept. 2022), pp. 886–894.
- [45] A. Bashirov and D. Balovnev. “Phantom Electroencephalogram Simulator for Automated Tests and Validation: Signals Repeatability”. In: *2021 International Russian Automation Conference (RusAutoCon)*. Sochi, Russian Federation: IEEE, Sept. 2021, pp. 660–664.
- [46] A. Anzolin, J. Toppi, M. Petti, F. Cincotti, and L. Astolfi. “SEED-G: Simulated EEG Data Generator for Testing Connectivity Algorithms”. In: *Sensors* 21.11 (May 2021), p. 3632.
- [47] L. R. Krol, J. Pawlitzki, F. Lotte, K. Gramann, and T. O. Zander. “SEREEGA: Simulating Event-Related EEG Activity”. In: *Journal of Neuroscience Methods* 309 (Nov. 2018), pp. 13–24.
- [48] M. Marsousi, J. Alirezaie, and K. Umapathy. “A Flexible Approach for Simulating Physiological Signals”. In: *Physiological Measurement* 34.6 (June 2013), pp. 695–712.

- [49] K. Karhinoja, A. Vasankari, J.-P. Sirkiä, A. Airola, D. Wong, and M. Kaisti. *Flexible framework for generating synthetic electrocardiograms and photoplethysmograms*. 2024. arXiv: 2408.16291 [cs.LG].
- [50] C. L. Petersen, M. Görges, R. Massey, G. A. Dumont, and J. M. Ansermino. “A Procedural Electroencephalogram Simulator for Evaluation of Anesthesia Monitors”. In: *Anesthesia & Analgesia* 123.5 (Nov. 2016), pp. 1136–1140.
- [51] D. Alvarez-Estevez and R. Rijsman. “Haaglanden Medisch Centrum Sleep Staging Database”. In: *PhysioNet* (2022).
- [52] D. Alvarez-Estevez and R. M. Rijsman. “Inter-Database Validation of a Deep Learning Approach for Automatic Sleep Scoring”. In: *PLOS ONE* 16.8 (Aug. 2021). Ed. by Y. Qiu, e0256111.
- [53] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley. “PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals”. In: *Circulation* 101.23 (June 2000).
- [54] F. Abtahi, F. Seoane, K. Lindecrantz, and N. Löfgren. “Elimination of ECG Artefacts in Foetal EEG Using Ensemble Average Subtraction and Wavelet Denoising Methods: A Simulation”. In: *XIII Mediterranean Conference on Medical and Biological Engineering and Computing 2013*. Ed. by L. M. Roa Romero. Vol. 41. Cham: Springer International Publishing, 2014, pp. 551–554.
- [55] L. Xu, E. Peri, R. Vullings, C. Rabotti, J. P. Van Dijk, and M. Mischi. “Comparative Review of the Algorithms for Removal of Electrocardiographic Interference from Trunk Electromyography”. In: *Sensors* 20.17 (Aug. 2020), p. 4890.
- [56] G. L. Wallstrom, R. E. Kass, A. Miller, J. F. Cohn, and N. A. Fox. “Automatic Correction of Ocular Artifacts in the EEG: A Comparison of Regression-Based and Component-Based Methods”. In: *International Journal of Psychophysiology* 53.2 (July 2004), pp. 105–119.
- [57] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent Component Analysis*. 1st ed. Wiley, May 2001.
- [58] J. T. Gwin, K. Gramann, S. Makeig, and D. P. Ferris. “Removal of Movement Artifact From High-Density EEG Recorded During Walking and Running”. In: *Journal of Neurophysiology* 103.6 (June 2010), pp. 3526–3534.