Non-Mathematical Musings on Information Theory and Networked Musical Practice

PHIL STONE

School of Veterinary Medicine, University of California, Davis, USA
Email: pkstone@ucdavis.edu

Claude Shannon's 1948 paper 'A Mathematical Theory of Communication' provided the essential foundation for the digital/information revolution that enables these very pixels to glow in meaningful patterns and permeates nearly every aspect of modern life. Information Theory, born fully grown from this paper, has been applied and mis-applied to a multitude of disciplines in the last 70-odd years, from quantum physics to psychology. Shannon himself famously decried those jumping on the 'scientific bandwagon' of Information Theory without sufficient mathematical rigour. Nevertheless, having a brief personal connection to Dr Shannon (and being extremely grateful for it), I will take the liberty of colouring some of my experience with computer network music with less-than-rigorous insights gained from his work.

1. INTRODUCTION

In 1981, fresh out of university, I secured the best summer job of my life. At graduation, a classmate named Peggy said her dad was a tinkerer who wanted to build a juggling robot and needed an assistant for the summer. Since I had some early experience with computer programming and electronics, she thought I might be a good fit for the position. Peggy's last name was Shannon, and her father was Claude Shannon.[1]

Six years later, I joined John Bischoff, Chris Brown, Scot Gresham-Lancaster, Tim Perkis and Mark Trayle to form 'The Hub', one of the earliest network music ensembles. Over the intervening three decades, the Hub has built a musical practice based on the sharing of information via digital interconnection.

Hub pieces, at their essence, are about information exchange. I will explore Hub musical practice and the implications of our musical experiments, and occasionally quote some Information Theory without any pretence of mathematical rigour.

[1]Though I had little appreciation at that point in my life for Shannon's impact on basically everything that I would go on to study, I jumped at the offer. We did not complete a juggling robot that summer (Dr Shannon built *that* in the following year), but we did construct an intricate motorised diorama of three juggling clowns, and I had the privilege of spending a summer working with a brilliant, kind, voraciously curious and inventive man. He was also an excellent juggler and taught me how to juggle clubs and ride a unicycle!

2. INFORMATION AND ENTROPY

> First Shannon had to eradicate 'meaning'; … 'The meaning of a message is generally irrelevant,' he proposed cheerfully. (Gleick 2011: 219)

For readers unfamiliar with the work of Claude Shannon, a little context is in order. Shannon's 1948 paper 'A Mathematical Theory of Communication' (republished as a book with commentary by Warren Weaver as *The Mathematical Theory of Communication* in 1949) is widely considered among the most important academic papers of all time – Google Scholar ranks it at no. 4 (Van Norden, Maher and Nuzzo: 2014). It was published within six months of the first working transistor prototype, which debuted just down the hall from Shannon's office at Bell Labs in Manhattan (Gleick 2011: 231). It is difficult to say which event was more crucial to the coming information age.

The paper proposed an elegant, formal mathematical basis for the measurement and analysis of communication, or more generally, information. It changed the fundamental nature of the word 'information' from a vague and colloquially defined 'know it when you see it' concept into a measurable quantity with derivable characteristics. In a paper with no less than 23 theorems, the following is only a small subset of the important concepts introduced:

- the 'bit', defined as the smallest quantum of information, representing two possible states;
- channel capacity: the bits per second a communications connection can handle *with accuracy*;
- 'information' is orthogonal to 'meaning' (see above);
- noise, defined as unintended, non-predictable distortion of an information source
- information = entropy = surprise.

For nearly all these now-fundamental concepts, Shannon's paper marked their very first published appearance.

The linkage between information, entropy and surprise is most intriguing but is, at first glance, rather non-intuitive. Shannon showed how the information contained in a message sent as a stream of discrete

Organised Sound **26**(3): 327–332
doi:10.1017/S1355771821000418

symbols can be measured as a summation function of the probabilities of every symbol's occurrence, affected by all preceding symbols. Stated in more general terms, the amount of information in a message (i.e., its propensity to *inform*) is equivalent to its level of unpredictability.

Thus, a message that is very predictable contains less information than one that is not predictable. A sequence such as 'ABCDEF' conveys little new information by the time the 'F' is reached, as the next symbol is very likely to be 'G'. At the other extreme, a message consisting of an entirely random sequence of symbols is maximally informative. This is the non-intuitive part, but only because 'information' is conventionally conflated with 'meaning'. At the very outset, Shannon's declaration of independence for information from meaning freed him to discover its synonymous relationship with unpredictability. A random sequence may be devoid of 'meaning', but it is full of information. Shannon called this measure of information-as-unpredictability 'entropy'. I will use the slightly more musically relevant terms 'surprise' (or occasionally 'liveness' or 'unpredictability') interchangeably with 'entropy' and 'information' throughout the rest of this article.

Let us now examine network music, especially as practised by the Hub, in the light of these concepts.

3. NETWORK MUSIC

> It is now possible to explain what one means by the capacity *C* of a noisy channel. It is, in fact, defined to be equal to the maximum rate (in bits per second) at which useful information … can be transmitted over the channel. (Shannon and Weaver 1949: 21)

For the purposes of this article, I will use the term 'network music' as shorthand for a method of live musical practice that, at a minimum, includes a usefully-high-speed digital connection between one[2] or more performers.

'Usefully-high-speed' here intentionally echoes Weaver's phrase 'useful information', to describe any communication sufficiently fast enough to accomplish desired musical goals. The Hub's first-generation network hardware, home-brewed from single-board microcomputers and custom wire-wrapped circuitry, worked at an effective rate of about 2,000 bits per second; nevertheless, this turned out to be quite musically useful.

The Hub's predecessors, the League of Automatic Music Composers, employed ad hoc connections between parallel ports:

> At other times we connected via the KIM's interrupt lines which enabled an instantaneous response as one player could 'interrupt' another player and send a burst of musical data which could be implemented by the receiving program immediately. (Brown and Bischoff 2002)

Even the earliest microcomputers, such as the tiny yet redoubtable KIM-1s employed by the League, were quite capable of musically 'useful' communication speed, that is, in the kilobits-per-second range.[3]

4. HUB'S PRACTICE OF NETWORK MUSIC

> Information is a measure of one's freedom of choice when one selects a message. (Shannon and Weaver 1949: 9)

The Hub's method of work has remained (with a few exceptions) consistent over the years and across various network and music-making technologies. One member writes up a 'spec' (short for 'specification') for a new piece, which is a text document outlining the information to be exchanged in the piece, how that information is to be generated and distributed through the network, and how each member should respond to it. Anything *not* delineated by the spec is left up to each performer's discretion. For example, while a given spec may describe how performers are to select the next pitch they will play, it may say little or nothing about what timbre or duration of sound is to be used or have any requirements for the timing of events. A different spec may only require an immediate response to a trigger event and say absolutely nothing about the pitch of that response (e.g., John Bischoff's 2005 piece *Tesla Sync*). Generally, what is not specified may be improvised.

There is a wide variety of degrees of control in Hub piece specs. Some are reductively simple and permissive; the ultra-compact spec for Scot Gresham-Lancaster's *Noosphere* (2006) states 'I am spewing "/yourname/scot/timbre n" where n is 0.0>n>1.0. Do with it what you will.'[4] On the other extreme, the spec to my 1987 piece *Borrowing and Stealing*, weighed in at a prolix five pages, spelling out a data format for exchanging melodic motifs, suggestions for what to do with that information and sometimes veering into editorial about the new-fangled network that we had just built. Nevertheless, both pieces allowed a high degree of performer autonomy.

[2]A one-musician network is possible (and useful!), employing feedback, that is, self-connection.

[3]I will mostly ignore the *telematic* capability of network music in this article. While the Hub's first concert explored this very thing, and while I grant that it can be a useful technique (especially for initial rehearsal of a piece), given a choice, our group prefers to make music *together* where we can hear, see and interact with each other – that is, in the same room.

[4]All Hub piece spec quotations have been culled from materials prepared for a collection of Hub scores published in (Brümmer 2021), unless otherwise indicated. Also see liner notes (The Hub 2008).

https://doi.org/10.1017/S1355771821000418 Published online by Cambridge University Press

Again, the general Hub principle is 'Anything not specified may be improvised.'

There is what has come to be another guiding principle of the Hub: the instrument/system a performer uses to make their sound is left completely up to them; the one requirement is that it be capable of meaningful participation in the Hub network. This applies not only to the choice and design of one's sound-making system but also to the computer hardware and software language used to code the Hub piece. In our hard-earned experience, this diversity of approaches contributes to a pleasing diversity in the music produced.

Developing Hub pieces from spec to performable works can be a gruelling process. Writing the software and developing the sonic material can take dozens of hours per piece. Our first rehearsal of a new piece is usually chaotic and often frustrating. Inevitably, unforeseen complications arise with even the simplest specs when confronted with the reality of the network. Massive tweaking of the spec may be required when the desired results do not materialise, are simply not achievable in reality, or turn out to be less than musically interesting. Finally, the stress of debugging one's faulty code while everyone else's seems to be working just fine (and making lots of noise) cannot be overstated.

Because it can be so time-consuming and difficult to write and debug a Hub piece, in the 1990s we briefly experimented with 'piece sharing'. At that moment in time, we all had (or had access to) computers that could run the music languages SuperCollider and Max. This allowed one member to design and implement a piece in one of these languages, and everybody else could just copy and run it on their respective computers.

While it was a more efficient mode of work, musically it was less satisfying. We realised that much of what is interesting in a Hub piece arises from the variety of approaches we each take to 'solving' a given Hub spec, and we abandoned piece sharing rather quickly.

4.1. Surprise!

In retrospect, uniformity of spec implementation eliminated an important source of unpredictability. The notion of 'surprise' comes up again and again in Hub discussions. Successful pieces are described as having the capacity to delight with unexpected or *emergent* behaviour (see *Waxlips*, later in this section). Over-specification or excessive demands for conformity are viewed with suspicion for their potential to damage or destroy this capacity.

Hub practice has been one of counterbalance against the regimentation and precision that computers engender and even encourage. It is all too possible to use a musical network in completely unsurprising, non-live ways. The first commercial implementation of musical networking, MIDI, was introduced in the early 1980s. It was designed by the burgeoning electronic music industry to allow one player (or computer) to control one or more electronic music devices.[5] Software MIDI sequencers became popular (I co-wrote one of them, for which I probably deserve a stint in musical purgatory), and a great deal of the popular music of the 1980s and beyond took on a characteristically 'synth-y' uniformity and often exhibited a deficit of 'liveness'. I suspect a finger of blame for that may be pointed in MIDI's direction.

It should be noted that an information-content/liveness continuum exists *within* network music practices. The mere use of a network does not intrinsically guarantee high information content. As an extreme example, shipping a pre-recorded performance around a digital network conveys very little information, other than the bit-for-bit description of the performance – which is highly compressible, not at all 'live', and contains no further potential for surprise after the first hearing. Also consider the previously described tightly sequenced MIDI network, in a one-to-many (conductor/orchestra) configuration; an 'authoritarian' configuration of this sort tends towards low entropy/information/surprise.

In this sense, Hub pieces themselves have various levels of entropy. Perhaps this might be better described in this context as 'liveness' – as a more musically familiar quantification of the range of possibilities at any given moment of a performance. Each moment of fully live music exists at a crux of near-infinite possibility; like a sporting event or game of chance, anything might happen next. However, while a more tightly specified piece might be expected to engender fewer alternatives at any given moment, in practice, the complexity inherent in an interconnected high-speed network supports no such simplistic deduction.

Some Hub pieces are so deterministic (on the surface) that they operate in a 'hands off' manner. An example of this is Mark Trayle's *Simple Degradation* (1987):

> One performer generates and processes a waveform, simulating the response of a plucked string. This waveform is then broadcast on the computer network, the other performers using it for amplitude modulation (loudness variation). The rate at which the waveform is played back by the performers is determined by the performer who generated the waveform. The performers are free to choose whatever timbres and pitches they wish. The waveform may only be used for amplitude modulation.

[5]The Hub actually (ab)used MIDI for their network in the 1990s, but Scot Gresham-Lancaster had to essentially reprogram an Opcode MIDI interface to allow its use as a multiplayer network.

> Pitch may only change after one complete cycle of the waveform. (Trayle quoted in Brown and Bischoff 2002)

In performance this spec results in a series of a dozen or so similar but individually quite distinct sharp crescendos, with instantaneously erratic volume changes that vary asynchronously between performers and the average of which gradually decays to silence. The piece ends after a pre-agreed number of 'plucks' of the algorithmic string. Very little input is required from each performer, except to occasionally change the sustained sound that is modulated by the received amplitude.

Hub pieces that require so little input from performers are exceptions rather than the rule, but defying expectations, *Simple Degradation* has a performance complexity that belies its deterministic spec, almost certainly due to the richness of its source material – the stochastically generated 'plucked string' data.

In fact, the liveness of a Hub piece is not necessarily derivable from the relative complexity or rigidity of its spec. Generally, a Hub composer seeks to create a fertile ground for the sowing of a chaotic harvest, although it is not always obvious where this might be hidden in the spec. A beautiful example of this is the terse yet 'tight' specification of Tim Perkis's *Waxlips* (1991):

> The piece is simple. Each player does essentially the same thing: take key-down midi messages in, transform them in a regular way that I'll specify below, play the new transformed note and send out a copy of it to somebody. I'll 'seed' the process at the beginning of the piece or section by sending out a few notes to start.
>
> The transformation can be anything you want, within these limitations: One note in, one note out. For every possible midi note and channel input combination (127 * 5 or 635 total) you define a unique transform to some other midi note and channel combination. Within any one performance of the piece this mapping is fixed: each time a particular note on a particular channel is received the same transformed output is sent. No random number changes, knob or slider or button adjustments, no algorithms which depend on previous states of your machine or previous input. A simple, fixed mapping.

While the rules of the piece are stringent, the observant reader will note that nowhere in this spec is any mention of 'timing', 'waiting' or 'delay'! The result is a gloriously chaotic feedback-driven murder of notes flying around the network at MIDI rate (31.25 Kbits per second). Any given interlude of *Waxlips* is as likely to wander noisily and seemingly randomly for long stretches of time as it is to fall into and then back out of fascinating patterns. The piece also served as a dependable 'coalmine canary' throughout its performance history, sensitive to any errors or other irregularities in the network (see 'noise', insection 4.3). Because of this it became the de-facto 'shakedown' piece for Hub tech rehearsals; if there were *any* 'holes' in a given network setup, from either faulty hardware, buggy software, or incorrect connection, the piece would not work correctly. Performances could vary wildly from setup to setup. All this rich and unpredictable behaviour arose from a spec that would seem, at first glance, like it might engender rather limited surprise! In this, Tim achieved his stated goal:

> I'd like to try to do one where you can really see if there is any emergent pattern to a static setup, where each station acts in a fixed, predictable way, but the interconnects are so complex that the overall behaviour is still groovy.

Members of the Hub share a consensus that surprise is to be encouraged and is a desirable characteristic of our pieces; in practice, this goal is reached via many avenues. Unpredictability is inherent in the independent implementation of pieces by each member and is encouraged by specifications that leave great latitude to each performer's improvisation and discretion. It is also true that the potential for unpredictability in a given Hub piece is not always obvious from its spec. Sometimes (if we do things right) unpredictable behaviour emerges from nothing more than the inherent complexity of a high-speed network, as *Waxlips* so perfectly illustrates.

Another technique often employed by the Hub to encourage surprise and liveness is important enough to deserve closer examination in the next section.

4.2. Unconventional musical power structures

Music ensembles tend to be organised under two broad categories. In the small ensemble or 'band' tradition, there is a collaborative, often democratic sharing of power. At the opposite end of the spectrum is the classical orchestra, with the conductor very much in charge. In both cases, the individual musicians control the moment-to-moment details of the sound they produce, that is, pitch, dynamics, timbre, timing, articulation and so on. Also, in both cases, particular instruments tend to have somewhat fixed roles in the ensemble. While the Hub very much gravitates towards the 'band' model, network interconnection gives us the ability to play with unconventional power distribution and role structures. The earliest piece to explicitly explore this was Tim Perkis's *Minister of Pitch* (1988):

> In most ensemble playing, the player of each instrument has some sort of responsibility for a particular aspect of the music: a bassline, keeping a rhythm, playing a main melody, etc. Using the capabilities offered when data is continuously exchanged between players, I organised responsibilities differently: one player's actions control the pitches played by all the players, another controls

https://doi.org/10.1017/S1355771821000418 Published online by Cambridge University Press

timing, and another is responsible for setting an overall timbre.

My piece *Boss* (2005) might well have been titled *Minister of Amplitude*, considering the debt it owes to Tim's groundbreaking re-apportioning of roles. But it injected a little anarchy into the mix (as it were):

> There is exactly one 'Boss' at any given time, who is completely in charge of the Hub's mix. The Boss sends amplitude messages to the other members of the group, in any way he or she desires. Anybody can become Boss at any time by deposing the former Boss. Sending a 'trigger' message to the HUB effects this coup d'amplitude. Upon receipt of such a message, the former boss must stop sending amplitude messages immediately. The new boss cranks out amplitude messages until similarly deposed. Non-Boss Hubsters may not make any adjustment to their own amplitude, except to tweak their maximum loudness to balance with the group.

Although *Boss* has some designed-in potential to foment a ruinous lack of cooperation (who has not felt the impulse to pull an ensemble-mate's fader down and turn one's own volume up?), a perhaps predictable result was that the group tended to use discretion (and occasional mischief) in taking over the 'Boss' role and mixing each other. The feeling of mixing a Hub piece live was quite enjoyable, though it could be challenging to share attention between the making of one's own sounds and controlling the group's mix at the same time.

Chris Brown's *Cut to Ribbons* (2006) approached the 'role-shuffling' technique of *Minister of Pitch* in a different way. In this piece, each player has eight sliders, controlling pitch, amplitude, timbre, tempo, rhythm, duration, density and phrasing. These parameters are used to compose (in real-time, i.e., algorithmically) 'phrase-streams', as Chris termed them. The catch was that although each performer could adjust their own sliders, so too could any other member of the group, at any time.

In all of these pieces, performers give up control over one or more parameters that musicians traditionally micromanage. At the same time, they are able to virtually 'reach into' their bandmates' systems and tweak their parameters. These three pieces produced dramatically different results, but all had in common a wonderful complexity and enormous capability for surprise. Upsetting the power structure/roles/apportionment of control of the ensemble is a fruitful source of 'good' entropy, and the power of digital interconnection makes it possible.

4.3. Bugs and other snafus: 'noise' and network music

In his commentary on Shannon's *Mathematical Theory of Communication*, Weaver writes:

> It is generally true that when there is noise, the received signal exhibits greater information – or better, the received signal is selected out of a more varied set than the transmitted signal. This is a situation which beautifully illustrates the semantic trap into which one can fall if he does not remember that 'information' is used here with a special meaning that measures freedom of choice and hence uncertainty as to what choice has been made. It is therefore possible for the word information to have either good or bad connotations. Uncertainty which arises by virtue of freedom of choice on the part of the sender is desirable uncertainty. Uncertainty which arises because of errors or because of the influence of noise is undesirable uncertainty. (Shannon and Weaver 1949: 19)

My preceding tongue-in-cheek allusion to 'good' entropy is a setup for Weaver's connoisseur-like discrimination between 'desirable' and 'undesirable' information. In one sense, he is entirely correct: there is a fragility to many Hub pieces, in that everybody must be working exactly correctly or else the entire thing fizzles (see *Waxlips*, section 4.1). In this sense, bugs are mostly characterised as very *undesirable* noise. An unexpected by-product of error or malfunction is the occasional production of some 'good' entropy, however. Again, *Waxlips* serves as an excellent example: even when the piece was not working exactly correctly, its behaviour could still be delightful, getting into strange cul-de-sacs of data that produced endlessly repeating loops sounding uncannily like composed motifs.

The point here is that the good information/bad information dichotomy may not be as clear-cut as Weaver seems to suggest (and to be fair to him, he is just trying to shake us free from our preconceived notion that *more* information is always better). Nevertheless, the Hub has found that 'desirable information' can occasionally arise even from unwanted noise and bugs.

5. CONCLUSION

The Hub has been making network music together more or less continuously since 1987. We lost our bandmate and friend Mark Trayle in 2015, and Matt Ingalls has joined us in our most recent work. The underlying technology has changed drastically over the years, but our method of work has remained consistent. We have seen the network hardware speed up, shrink, and finally, disappear (Bencina n.d.). There may be some nostalgia for the old home-brewed networks, microcomputers and sound-making gear – tables overflowing with horrendous tangles of interconnecting cables – but none of us would want to go back to lugging piles of anvil cases to gigs that take hours to set up and troubleshoot. Fast laptops capable of real-time audio synthesis have been a blessing to us network musicians of a certain age: they allow us to

https://doi.org/10.1017/S1355771821000418 Published online by Cambridge University Press

explore the things that we love about this type of music without pulling muscles. It should be noted that though we may have settled into more homogenous and compact hardware, each member still writes unique 'solutions' to Hub specs, in a variety of programming and synthesis environments.

The label of 'network music' has come to be applied to many other practices besides ours; in the shadow of the recent pandemic, musicians (and fans of live music) have figured out ways to interact over the internet, in near-real-time or (mostly) otherwise. It would be gatekeeping to insist that this is not network music; clearly it is, in the most *literal* sense. In terms suggested earlier, however, this is *low-entropy* network music; it serves a worthy purpose of making musical interaction even possible in a time of isolation, but it is still a poor substitute for actual co-located live music. In the post-pandemic world, 'internet' music will remain useful for rehearsals and collaborations that cannot be accomplished in-person. It is hoped, however, that 'network music' will not come to be synonymous with music that merely happens to travel over the internet.

Speaking of 'horrendous tangles of interconnecting cables', there is a current revival of interest in analogue modular synthesisers, and it is thrilling to see and hear this (at least occasionally) as a live performance practice. Electronic (whether analogue or digital) improvisation, in the tradition of MEV and many others, has been a core part of the Hub since day one – it is the chaotic base layer upon which we project the structure of our information exchange specs. It is encouraging to see a new chapter of live electronic music unfolding. Perhaps a hybrid practice, an analogue/network band, might prove rewarding for the strong-of-back and logical-of-mind.

I mean by this to show that a 'network band' can take many forms, as long as there is a usefully high-speed digital interconnection involved. This is trivial to achieve these days (though the writing of the necessary code is admittedly less trivial). The Hub's experience would suggest a useful principle for such bands in whatever form they might assume: cherish entropy! We all live under the inescapable oppression of the Second Law of Thermodynamics – life can sometimes seem like a Sisyphean struggle against the relentless creep of disorder. Claude Shannon gifted us with a different way to view entropy, as being full of information. How wonderful it is that there is astonishment and beauty to be found there!

REFERENCES

Bencina, R. n.d. OSC Groups. www.rossbencina.com/code/oscgroups (accessed 15 March 2021).

Brown, C. and Bischoff, J. 2002. Indigenous to the Net: Early Network Music Bands in the San Francisco Bay Area. http://crossfade.walkerart.org/brownbischoff/IndigenoustotheNetPrint.html (accessed 24 February 2021).

Brümmer, L. (ed.) 2021. *The Hub: Pioneers of Network Music*. Heidelberg: Kehrer Verlag.

Gleick, J. 2011. *The Information: A History, a Theory, a Flood.* New York: Vintage Books.

The Hub. 2008. Liner notes to *Boundary Layer*. 3 CDs. New York: Tzadik, TZ 8050-3.

Shannon, C. E. and Weaver, W. 1949. *The Mathematical Theory of Communication*. Urbana and Chicago: University of Illinois Press.

Van Norden, R., Maher, B. and Nuzzo, R., 2014. The Top 100 Papers. *Nature*. www.nature.com/news/the-top-100-papers-1.16224#/alternative (accessed 24 February 2021).