

# The Free Energy of a General Computation

Ämin Baumeler\* and Stefan Wolf†

\* Institute for Quantum Optics and Quantum Information (IQOQI), Austrian Academy of Sciences, Boltzmannngasse 3, 1090 Vienna, Austria

† Faculty of Informatics, Università della Svizzera italiana, 6900 Lugano, Switzerland

**Abstract**—Starting from Landauer’s slogan “information is physical,” we revise and modify Landauer’s principle stating that the *erasure* of information has a minimal price in the form of a certain quantity of free energy. We establish a direct link between the *erasure cost* and the *work value* of a piece of information, and show that the former is essentially the length of the string’s best compression by a reversible computation. We generalize the principle by deriving bounds on the free energy to be invested for — or gained from, for that matter — a *general computation*. We then revisit the *second law of thermodynamics* and compactly rephrase it (assuming the Church/Turing/Deutsch hypothesis that physical reality can be simulated by a universal Turing machine): Time evolutions are *logically reversible* — “the future fully remembers the past (but not necessarily *vice versa*).” We link this view to previous formulations of the second law, and we argue that it has a particular feature that suggests its “logico-informational” nature, namely *simulation resilience*: If a computation faithfully simulates a physical process violating the law — then that very computation procedure violates it as well.

## I. INTRODUCTION

In 1961, Rolf Landauer famously stated “Information is physical” [13]: Despite the success of Shannon’s making information an *abstract* concept (that can be viewed and understood independently of its particular physical realization), Landauer — while not questioning the power of that abstract view — recalls that all information storing, treatment, and transmission is ultimately a *physical* process and, thus, subject to physical laws. A specific law relevant in this context is the second law of thermodynamics. Its consequence for information processing has been called *Landauer’s principle* [12]: “The erasure of  $N$  bits of information costs at least an amount of  $NkT \ln 2$  ( $k$  being Boltzmann’s constant) of free energy that must be dissipated as heat into the environment of temperature  $T$ .” (Note that this heat dissipation is actually crucial for the argument: It represents the compensation required for avoiding a violation of the second law despite an entropy decrease in the memory device through the erasure process.)

In this note, we modify and generalize Landauer’s principle in the following respects: First, it is claimed that the erasure cost is not proportional to the length of the string to be erased, but of its best *compression* — given the entire knowledge of the erasure device (Section IV). We obtain these results from new bounds on the *work value* of information (Section III), and a direct connection between erasure cost and work value of any piece of information. Second, we generalize these results to a lower bound on the free-energy cost, or value, of a *general* computation (Section V). Our findings are modifications of

known results (e.g., [7], [9]) to the *constructive* setting — where all involved processes are imagined to be carried out by a Turing machine. Finally, we also give a lower bound on the free-energy *gain* possible from certain computations — a bound that matches the *cost* of the inverse computation. We look at the use of the erasure cost as an *intrinsic randomness* measure in the context of quantum correlations (Section VI). We finish by proposing a computational version of the second law of thermodynamics (Section VII). This comes with a speculation about what trait of it is the reason why such a version exists in the first place. Candidates are its “encoding independence” and “simulation resilience.” If a computation simulates a process violating the second law, then that computing procedure *cannot be closed* but must dissipate “junk” bits onto the other parts of the tape, or heat into the environment: The violation of the law by a process carries over to *its simulation*. The reason is that a degree of freedom is represented by — a degree of freedom.

## II. WORK VALUE: STATE OF THE ART

Bennett [3] claimed the work value of a string  $S$  to be proportional to the difference between its length and the length of the shortest program that produces  $S$ . The latter is called the *Kolmogorov complexity*  $K(S)$  of  $S$  [11]:

$$\text{WV}(S) = (\text{len}(S) - K(S))kT \ln 2 .$$

Bennett’s argument is that  $S$  can be logically, hence, thermodynamically [10] *reversibly* mapped to the string  $P||000 \cdots 0$ , where  $P$  is the shortest program generating  $S$ , and the length of the generated 0-string is  $\text{len}(S) - K(S)$ .

It was already pointed out by Zurek [20] that while it is true that the reverse direction exists and is computable by a universal Turing machine, its *forward direction*, i.e., obtaining  $P$  from  $S$ , is *uncomputable*. This means that a “demon” that could carry out this work-extraction computation on  $S$  *does not exist* (if the Church-Turing hypothesis is true). We will see, however, that Bennett’s value is an *upper bound* on the work value of  $S$ . Bennett also links the string’s erasure cost to its probabilistic entropy [5].<sup>1</sup>

<sup>1</sup>Dahlsten *et al.* [7] follow Szilárd [14] in putting the *knowledge* of the demon extracting the work to the center of their attention. More precisely, they claim  $\text{WV}(S) = \text{len}(S) - D(S)$ , where the “defect”  $D(S)$  is bounded from above and below by a *smooth Rényi entropy* of the distribution of  $S$  from the demon’s viewpoint, modeling its ignorance. Building on these results and in the same probabilistic spirit, the cost of erasure [8] as well as of general computations [9] have been linked to entropic expressions of (conditional) probability distributions.

### III. WORK EXTRACTION AS DATA COMPRESSION

In the following, we model work extraction to be an algorithm executed by a “demon with knowledge.”

#### A. The Model

We assume the demon to be a *universal Turing machine*  $\mathcal{U}$  the memory tape of which is sufficiently long for the inputs and tasks in question, but *finite*. The tape initially contains  $S$ , the string the work value of which is to be determined,  $X$ , a finite string modeling the demon’s *knowledge about*  $S$ , and 0s for the rest of the tape. After the extraction computation, the tape contains, at the bit positions initially holding  $S$ , a (shorter) string  $P$  plus  $0^{\text{len}(S)-\text{len}(P)}$ , whereas the rest of the tape is (again) the same as before the work extraction. The operations are *logically* reversible and can, hence, be carried out *thermodynamically* reversibly [10]. Logical reversibility is the ability of the same demon to carry out the backward computation step by step, *i.e.*, from  $P||X$  to  $S||X$ . We denote by  $\text{WV}(S|X)$  the *maximal length of an all-0-string extractable logically reversibly from*  $S$ , *given the knowledge*  $X$ , times  $kT \ln 2$ , *i.e.*,

$$\text{WV}(S|X) := (\text{len}(S) - \text{len}(P))kT \ln 2$$

if  $P$ ’s length is minimal.

#### B. Lower Bound

We show that every specific data-compression algorithm leads to a lower bound on the extractable work: Let  $C$  be a computable function

$$C : \{0, 1\}^* \times \{0, 1\}^* \longrightarrow \{0, 1\}^*$$

such that

$$(A, B) \mapsto (C(A, B), B)$$

is injective. We call  $C$  a *data-compression algorithm with helper*. Then we have

$$\text{WV}(S|X) \geq (\text{len}(S) - \text{len}(C(S, X)))kT \ln 2.$$

This can be seen as follows. First, note that the function

$$A||B \mapsto C(A, B)||0^{\text{len}(A)-\text{len}(C(A, B))}||B$$

is computable and bijective. From the two (possibly irreversible) circuits which compute the compression and its inverse, one can obtain a *reversible* circuit for the function such that no further input or output bits are involved: This can be achieved by first implementing all logical operations with Toffoli gates and uncomputing the “junk” [4] in both circuits. The resulting two circuits have now still the property that the input is part of the output. As a second step, we can simply combine the two such that the first circuit’s first and second outputs become the second’s second and first inputs, respectively. Roughly speaking, the first computes the compression and the second reversibly uncomputes the raw data. The combined circuit has only *the compressed data plus the 0s* as the output, sitting on the bit positions carrying the input before. (This circuit is roughly as efficient as the less efficient of the two irreversible circuits for data compression

and decompression, respectively.) A typical example for an algorithm that can be used here is universal data compression *à la* Ziv-Lempel [19].

#### C. Upper Bound

We have the following upper bound on the extractable work:

$$\text{WV}(S|X) \leq (\text{len}(S) - K_{\mathcal{U}}(S|X))kT \ln 2,$$

where  $K_{\mathcal{U}}(S|X)$  is the conditional Kolmogorov complexity (with respect to the universal Turing machine  $\mathcal{U}$ ) of  $S$  given  $X$ , *i.e.*, the length of the shortest program  $P$  for  $\mathcal{U}$  that outputs  $S$ , given  $X$ . The reason is that the extraction demon is only able to carry out the computation in question (logically, hence, thermodynamically) reversibly if it is able to carry out the reverse computation as well. Therefore, the string  $P$  must be at least as long as the shortest program for  $\mathcal{U}$  generating  $S$  if  $X$  is given.

Although the same is not true in general, this upper bound is *tight* if  $K_{\mathcal{U}}(S|X) \approx 0$ . The latter means that  $X$  itself can be seen as a program for generating an additional copy of  $S$ . The demon can then bit-wisely XOR this extra  $S$  to the original  $S$  (to be work-extracted) on the tape, hereby producing  $0^{\text{len}(S)}$  *reversibly* to replace the original  $S$ , at the same time saving the new one, as reversibility demands. When Bennett’s “uncomputing trick” is used — allowing to make any computation by a Turing machine logically reversible [4] —, then a history string  $H$  is written to the tape during the computation of  $S$  from  $X$  such that after XORing, the demon can, in a (reverse) stepwise manner, *uncompute* the generated copy of  $S$  and end up in the tape’s original state — except that the original  $S$  is now replaced by  $0^{\text{len}(S)}$ : This results in a maximal work value matching the (in that case trivial) upper bound.<sup>2</sup>

### IV. REVISING LANDAUER’S PRINCIPLE

#### A. Connection to Work Value

For a string  $S \in \{0, 1\}^N$ , let  $\text{WV}(S|X)$  and  $\text{EC}(S|X)$  be its work value and erasure costs, respectively, given an additional string  $X$  (a “catalyst” which remains unchanged, as above). Then

$$\text{WV}(S|X) + \text{EC}(S|X) = NkT \ln 2.$$

<sup>2</sup>Let us compare our bounds with the entropy-based results of [7]: According to the latter, a demon *knowing*  $S$  *entirely* is able to extract maximal work:  $\text{WV}(S) \approx \text{len}(S)$ . What does it mean to “know  $S$ ”? The knowledge can consist of (a) a *copy* of  $S$ , or of (b) the ability to *compute* such a copy with a given program  $P$ , or (c) it can determine  $S$  uniquely *without* providing the ability to compute it. The constructive as opposed to the entropic groups of results are in accordance in the cases (a) and (b) but *in conflict* in case (c): For instance, assume the demon’s knowledge about  $S$  is: “ $S$  equals the first  $N$  bits  $\Omega_N$  of the binary expansion of  $\Omega$ ,” where,  $\Omega$  is the so-called halting probability [6] of a fixed universal Turing machine  $\mathcal{A}$  (*e.g.*, the demon  $\mathcal{U}$  itself). Although there is a short *description* of  $S$  in this case, and  $S$  is thus uniquely determined in an *entropic* sense, it is still incompressible, even given that knowledge:  $K_{\mathcal{U}}(\Omega_n | \text{“It is bits } 1-n \text{ of TM } \mathcal{A}\text{’s halting probability”}) \approx n$ : No work is extractable according to our upper bound. Intuitively, this gap opens up whenever the “*description complexity*” is smaller than the *Kolmogorov complexity*. (Note that a self-reference argument, called *Berry paradox*, shows that the notion of “description complexity” is problematic and can never be defined consistently for all strings.)

To see this, consider first the combination extract-then-erase. Since this is *one specific way* of erasing, we have

$$\text{EC}(S|X) \leq NkT \ln 2 - \text{WV}(S|X) .$$

If, on the other hand, we consider the combination erase-then-extract, this leads to

$$\text{WV}(S|X) \geq NkT \ln 2 - \text{EC}(S|X) .$$

### B. Bounds on the Erasure Cost

Given the results on the work value above, as well as the connection between the work value and erasure cost, we obtain the following bounds on the thermodynamic cost of erasing a string  $S$  by a demon, modeled as a universal Turing machine  $\mathcal{U}$  with initial tape content  $X$ .

**Landauer's principle, revisited.** *Let  $C$  be a computable function  $C : \{0,1\}^* \times \{0,1\}^* \rightarrow \{0,1\}^*$  such that  $(A, B) \mapsto (C(A, B), B)$  is injective. Then we have*

$$K_{\mathcal{U}}(S|X)kT \ln 2 \leq \text{EC}(S|X) \leq \text{len}(C(S, X))kT \ln 2 .$$

## V. GENERALIZING LANDAUER'S PRINCIPLE

Assume that a (universal) Turing machine performs a computation such that the initial content of the tape is  $A$  and  $X$  (plus a corresponding finite number of 0s) and the final state is  $B$  and  $X$  (where  $X$  can be seen, again, as a ‘‘catalyst’’). Depending on  $A$ ,  $B$ , and  $X$ , this computation can have a work cost or value, respectively.

### A. The Energy Cost of a General Computation

The following result is an algorithmically constructive modification of entropic results [9] and a generalization of less constructive but also complexity-based claims [21].

**Work cost of a general computation.** *Let  $C$  be a computable function,  $C : \{0,1\}^* \times \{0,1\}^* \rightarrow \{0,1\}^*$ , such that  $(V, W) \mapsto (C(V, W), W)$  is injective. Assume that the Turing machine  $\mathcal{U}$  carries out a computation such that  $A$  is its initial state,  $C_1$  the first intermediate state,  $C_2$  the second, etc., up to  $C_n$ , and  $B$  is the final state. Then the energy cost of this computation with side information  $X$  (always on the tape),  $\text{Cost}_{\mathcal{U}}(A \rightarrow_{\{C_i\}} B | X)$ , is at least*

$$\begin{aligned} \text{Cost}_{\mathcal{U}}(A \rightarrow_{\{C_i\}} B | X) &\geq kT \ln 2 \cdot [K_{\mathcal{U}}(A|X)] \dots \\ &- \sum_{i=1}^n (\text{len}(C(C_i|X)) - K_{\mathcal{U}}(C_i|X) - \text{len}(C(B, X))) . \end{aligned}$$


*Proof.* Let us consider the computation from  $(A, X)$  to  $(C_1, X)$ . According to the above, the erasure cost of  $A$ , given  $X$ , is at least  $K_{\mathcal{U}}(A|X) \cdot kT \ln 2$ . One possibility of realizing this complete erasure of  $A$  is to first transform it to  $C_1$  (given  $X$ ), and then erase  $C_1$  — at cost at most  $\text{len}(C(C_1, X)) \cdot kT \ln 2$ . Therefore, the cost to get from  $A$

to  $C_1$  given  $X$  cannot be lower than the difference between  $K_{\mathcal{U}}(A|X) \cdot kT \ln 2$  and  $\text{len}(C(C_1, X)) \cdot kT \ln 2$ . The statement follows by summing all contributions of the individual computing steps. *qed.*

Note that if no intermediate results are specified, the bound simplifies to

$$\text{Cost}_{\mathcal{U}}(A \rightarrow B | X) \geq [K_{\mathcal{U}}(A|X) - \text{len}(C(B, X))] \cdot kT \ln 2$$

(see also [18]).



$$\text{Cost}(A \rightarrow B | X) \geq [K(A|X) - \text{len}(C(B, X))] \cdot kT \ln 2$$

Figure 1. The energy cost of a general computation.

### B. The Energy Value of a General Computation

We consider the *work value of a computation* from  $A$  to  $B$ , given  $X$ . More specifically, this is a computation that starts with  $(A, X)$  and finishes with  $(B, X)$ , where  $B$  is *freely choosable* by the computation among all strings with a given complexity  $K_{\mathcal{U}}(B|X)$ . The work value is denoted by  $\text{WV}_{\mathcal{U}}(A \rightarrow B | X)$ , and it is bounded from below as follows.

**Work value of a general computation.** *Let  $C$  be a computable function,  $C : \{0,1\}^* \times \{0,1\}^* \rightarrow \{0,1\}^*$ , such that  $(V, W) \mapsto (C(V, W), W)$  is injective. The work value of a computation from  $A$  to  $B$ , given  $X$ , is bounded from below by*

$$\text{WV}_{\mathcal{U}}(A \rightarrow B | X) \geq [K_{\mathcal{U}}(B|X) - \text{len}(C(A, X))] \cdot kT \ln 2 .$$

*Proof.* The cost of erasing  $A$ , given  $X$ , is at most  $\text{len}(C(A, X)) \cdot kT \ln 2$ . We use a stretch of the resulting all-0-string of some length  $N$  for gaining  $NkT \ln 2$  free energy. The resulting string of length  $N$  is then used as a program for the universal Turing machine  $\mathcal{U}$ , with additional input  $X$ , and where the computation is made logically reversible using Bennett's ‘‘uncomputing’’ trick [4]; let  $B$  be the resulting string. Then  $K_{\mathcal{U}}(B|X) \leq N$ . *qed.*

### C. Combination

Consider the following “circular computation,” given  $X$ :

$$A \longrightarrow B \longrightarrow A .$$

The free-energy gain of computing  $B$  from  $A$  is at least

$$[K_{\mathcal{U}}(B|X) - \ln(C(A, X))] \cdot kT \ln 2 ,$$

whereas the cost for computing  $A$  back from  $B$  is at least this same amount. The identity of the two bounds is not very surprising; it implies that the bound on the work value is “at least as tight” as the one for the cost of the inverse computation, since otherwise a *perpetuum mobile of the second kind* results.

### VI. “RANDOMNESS” AND QUANTUM CORRELATIONS

Landauer’s revised principle suggests that the erasure cost of a piece of information is an *intrinsic, context-free, physical measure for its randomness* independent of probabilities and counterfactual statements (that “some value *could* just as well have been *different*”). This can be tested in a context in which randomness is central: *Bell correlations* [2] predicted by quantum theory. In a proof of principle, it was shown [16] that in essence, a similar mechanism as in the probabilistic setting arises: *If the correlation is non-local, the inputs are incompressible, and non-signaling holds, then the outputs must be highly complex as well.* This allows for a discussion of quantum correlations without the usual counterfactual arguments used in derivations of *Bell inequalities* (combining in a single formula results of different measurements that cannot actually be carried out together). Furthermore, this potentially opens the door to novel functionalities, namely *complexity amplification and expansion* [1]. What results is an *all-or-nothing flavor of the Church-Turing hypothesis* [17]: Either no physical computer exists that is able to produce non-Turing-computable data — or even a “device” as simple as a single photon can.

### VII. THE SECOND LAW AS LOGICAL REVERSIBILITY

In Landauer’s principle, the price for the *logical irreversibility* of the erasure transformation comes in the form of a *thermodynamic effort*. (Since the amount of the required free energy, and heat dissipation, is proportional to the length of the best compression of the string, the latter can be seen as a *quantification* of the erasure transformation’s irreversibility.) In an attempt to harmonize this somewhat *hybrid* picture, we invoke Wheeler’s [15] “*It from Bit*: Every *it* — every particle, every field of force, even the spacetime continuum itself — derives its function, its meaning, its very existence entirely [...] from the apparatus-elicited answers to yes-or-no questions, binary choices, *bits*.” This is an anti-thesis to Landauer’s slogan, and we propose the following *synthesis* of the two: If Wheeler motivates us to look at the environment as being a *computation* as well, then Landauer’s principle may be read as: The necessary environmental compensation for the logical irreversibility of the erasure of  $S$  is such

that *the overall computation, including the environment, is logically reversible: no information ever gets completely lost.*

### Second law, logico-computational version.

*Time evolutions of closed systems are injective:*

*Nature computes with Toffoli, but no AND or OR gates.*

Note that this fact is *a priori* asymmetric in time: The future must uniquely determine the past, not necessarily *vice versa*. In case the condition holds also for the reverse time direction, the computation is called *deterministic*, and *randomized* otherwise.

Logical reversibility is a simple computational version of a discretized second law, and it has implications resembling the traditional versions of the law: First of all, it leads to a “Boltzmann-like” form, *i.e.*, the existence of a quantity essentially monotonic in time. More specifically, the logical reversibility of time evolution implies that the Kolmogorov complexity of the global state at time  $t$  can be smaller than the one at time 0 only by at most  $K(R_t) + O(1)$  if  $R_t$  is a string encoding the time span  $t$ . The reason is that one possibility of describing the state at time 0 is to give the state at time  $t$ , plus  $t$  itself; the rest is exhaustive search using only a constant-length program simulating forward time evolution (including possible randomness).

Similarly, logical reversibility also implies statements resembling the version of the second law due to *Clausius*: “Heat does not spontaneously flow from cold to hot.” The rationale here is explained with a toy example: If we have a circuit — the time evolution — using only (logically reversible) Toffoli gates, then it is *impossible* that this circuit computes a transformation mapping a pair of strings to another pair such that the Hamming-weight of the two becomes even heavier while the lighter gets lighter. A function *accentuating imbalance*, instead of lessening it, is not injective, as the following counting argument shows.

“*Clausius*” Toy Example. Let a circuit consisting of only Toffoli gates map an  $N(= 2n)$ -bit string to another. We consider the map separately on the first and second halves and assume the computed function to be conservative, *i.e.*, to leave the Hamming weight of the full string unchanged at  $n$  (conservativity can be seen as some kind of *first law*, *i.e.*, the preservation of a quantity). We look at the excess of 1s in one of the halves (which equals the deficit of 1s in the other). We observe that the probability (with respect to the uniform distribution over all strings of some Hamming-weight couple  $(wn, (1-w)n)$ ) of the *imbalance substantially growing* is exponentially weak. The key ingredient for the argument is the function’s injectivity. Explicitly, the probability that the weight couple goes from  $(wn, (1-w)n)$  to  $((w + \Delta)n, (1-w - \Delta)n)$  — or more extremely —, for

$1/2 \leq w < 1$  and  $0 < \Delta \leq 1 - w$ , is

$$\frac{\binom{n}{(w+\Delta)n} \binom{n}{(1-w-\Delta)n}}{\binom{n}{wn} \binom{n}{(1-w)n}} = 2^{-\Theta(n)}.$$

The example suggests that logical reversibility might be the “Church/Turing manifestation” of the second law: If reality is computed by a Turing machine, then physical laws correspond to properties of such computations — as in the case of the second law: logical reversibility. If we assume for a moment that the second law of thermodynamics has indeed such a simple Church/Turing manifestation, it is a natural question in how far just this already makes the law special. In fact, the law *does* have a peculiar related property, *encoding independence*: Since the second law deals with degrees of freedom, and a degree of freedom will correspond, in another encoding, to a degree of freedom again, it is either respected in both encodings or violated in both. (In comparison: It cannot be decided by just looking at a running program whether the simulated system “violates or respects Kepler’s laws” — that would crucially depend on how masses and their position are represented by the code.) Hand in hand with this comes the property of *simulation resilience*. Let us take again the example of Kepler’s laws: If a program simulates planets moving on, say, *square* orbits, this does not mean that the program execution, viewed as a physical process, is by itself problematic — the laws of gravity are *not* simulation resilient. If, in sharp contrast to this, we simulate a system, *e.g.*, the microstate sequence of a steam engine, then that simulation process violates the second law just as the simulated system does: *With respect to the second law of thermodynamics, the simulation of “reality” is just as good as “reality” itself.*

## VIII. CONCLUSION

We start from Landauer’s principle, stating that the erasure of information requires an amount of free energy, to be dissipated as heat into the environment, proportional to the number of independent binary degrees of freedom of that information. Specifically, the use of reversible data compression, imagined to be carried out by Fredkin and Toffoli’s ballistic computer, implies that the necessary amount is proportional to the length of the best compression of the information into a binary string (and not to the length of the original string, as often stated). We generalize and broaden the scope of the principle, and its converse, to lower bounds on the free-energy cost of — or gain from — a *general computation*: the bounds on cost *versus* gain are in accord.

Landauer has derived, in 1961, his principle from the second law of thermodynamics. We close the circle by formulating a simple “Church/Turing version” of that law: *Logical reversibility of the overall computation, including the environment.* This fact alone implies variants of the historical versions of the second law, due to Boltzmann, Clausius, also Kelvin; it is perhaps equivalent to them, certainly simpler. The arising belief that the law is rather *logical* than *physical* in its nature

is nourished by two properties of the second law: *its encoding independence* and *its simulation resilience*.

Confronted with the relevance of the second law of thermodynamics in computation, and with its *simulation resilience*, let us close with the (provocative) question whether Landauer’s “*Information is physical*” should be replaced by

**“The second law of thermodynamics is not physical.”**

**Acknowledgments.** We thank Charles Bédard, Gilles Brassard, Claus Beisbart, Paul Boes, Časlav Brukner, Xavier Coiteux-Roy, Bora Dakic, Paul Erker, Jürg Fröhlich, and Arne Hansen for enlightening discussions. Ä.B. is supported by the Swiss National Science Foundation (SNSF) under grant 175860, and the Erwin Schrödinger Center for Quantum Science & Technology (ESQ). S.W. is supported by the Swiss National Science Foundation (SNSF), the NCCR *QSIT*, and the *Hasler Foundation*.

## REFERENCES

- [1] Ä. Baumeler, C. Bédard, G. Brassard, and S. Wolf, Kolmogorov amplification from Bell correlation, *International Symposium in Information Theory (ISIT) 2017*, 2017.
- [2] J. S. Bell, On the Einstein-Podolsky-Rosen paradox, *Physics*, Vol. 1, pp. 195–200, 1964.
- [3] C. H. Bennett, The thermodynamics of computation, *International Journal of Theoretical Physics*, Vol. 21, No. 12, pp. 905–940, 1982.
- [4] C. H. Bennett, Logical reversibility of computation, *IBM J. Res. Develop.*, Vol. 17, No. 6, pp. 525–532, 1982.
- [5] C. H. Bennett, Notes on Landauer’s principle, reversible computation and Maxwell’s demon, *Studies in History and Philosophy of Modern Physics*, Vol. 34, pp. 501–510, 2003.
- [6] G. Chaitin, A theory of program size formally identical to information theory, *Journal of the ACM*, Vol. 22, pp. 329–340, 1975.
- [7] O. Dahlsten, R. Renner, E. Rieper, and V. Vedral, The work value of information, *New J. Phys.*, Vol. 13, 2011.
- [8] L. del Rio, J. Åberg, R. Renner, O. Dahlsten, and V. Vedral, The thermodynamic meaning of negative entropy, *Nature*, 2011.
- [9] P. Faist, F. Dupuis, J. Oppenheim, and R. Renner, The minimal work cost of information processing, *Nature Communications*, 2015.
- [10] E. Fredkin and T. Toffoli, Conservative logic, *International Journal of Theoretical Physics*, Vol. 21, No. 3–4, pp. 219–253, 1982.
- [11] A. Kolmogorov, Three approaches to the quantitative definition of information, *Problem. Per. Informat.*, Vol. 1, No. 1, pp. 3–11, 1965.
- [12] R. Landauer, Irreversibility and heat generation in the computing process, *IBM Journ. of Res. and Developm.*, Vol. 5, pp. 183–191, 1961.
- [13] R. Landauer, Information is inevitably physical, *Feynm. & Comp.*, 1998.
- [14] L. Szilárd, Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen (On the reduction of entropy in a thermodynamic system by the intervention of intelligent beings), *Zeitschrift für Physik*, Vol. 53, pp. 840–856, 1929.
- [15] J. A. Wheeler, Information, physics, quantum: the search for link, *Proceedings III International Symposium on Foundations of Quantum Mechanics*, pp. 354–368, 1989.
- [16] S. Wolf, Non-locality without counterfactual reasoning, *Phys. Rev. A*, Vol. 92, No. 052102, 2015.
- [17] S. Wolf, An all-or-nothing flavor to the Church-Turing hypothesis, *Proceedings of Theory and Applications of Models of Computation (TAMC) 2017*, LNCS, Vol. 10185, pp. 39–56, Springer-Verlag, 2017.
- [18] S. Wolf, Second thoughts on the second law, in *Adventures Between Lower Bounds and Higher Altitudes*, LNCS, Vol. 11011, 2018.
- [19] J. Ziv and A. Lempel, Compression of individual sequences via variable-rate coding, *IEEE Transactions on Information Theory*, Vol. 24, No. 5, p. 530, 1978.
- [20] W. H. Zurek, Algorithmic randomness and physical entropy, *Phys. Rev. A*, Vol. 40, No. 8, 1989.
- [21] W. H. Zurek, Thermodynamic cost of computation, algorithmic complexity and the information metric, *Nature*, Vol. 341, pp. 119–124, 1989.