Out of equilibrium learning dynamics in physical allosteric resistor networks

Menachem Stern[*]
Department of Physics and Astronomy
University of Pennsylvania

Sam Dillavou
Department of Physics and Astronomy
University of Pennsylvania

Marc Z. Miskin
Department of Electrical and Systems Engineering
University of Pennsylvania

Douglas J. Durian
Department of Physics and Astronomy
University of Pennsylvania

Andrea J. Liu
Department of Physics and Astronomy
University of Pennsylvania

Abstract

Physical networks can learn desirable functions using local learning rules in space and time. Real learning systems, like natural neural networks, can learn out of equilibrium, on timescales comparable to their physical relaxation. Here we study coupled learning, a framework that supports learning in equilibrium in diverse physical systems. Relaxing the equilibrium assumption, we study experimentally and theoretically how physical resistor networks learn allosteric functions far from equilibrium. We show how fast learning produces oscillatory dynamics beyond a critical threshold, and that learning succeeds well beyond that threshold. These findings show how coupled learning may train systems much faster than assumed before, suggesting their applicability to slowly relaxing physical systems.

1 Introduction

Recent years have seen the explosion of learning based methods in the design and manufacturing of physical networks with desired properties [1, 2, 3, 4, 5], most of which use computational machine learning. More recently, new learning methods emerged, where the physical network *itself* is trained, evolving its parameters in order to perform desired tasks. Such physical learning machines have either minimized the energy of their desired response (directed aging) [6, 7, 8], or used contrastive learning [9, 10, 11, 12, 13] to minimize cost functions of interest. As an example, a network of variable resistors whose resistances are tuned by minimizing the cost function has “physical" degrees of freedom (e.g. currents on edges) and “learning" degrees of freedom (e.g. resistances of edges). The physical degrees of freedom adjust rapidly so that the system is always fully equilibrated with respect to the physical degrees of freedom at each step of the learning process.

Algorithms that rely on equilibration of physical degrees of freedom during learning have not seen wide use in machine learning [14] because the equilibration slows their performance. Recent work has relaxed that requirement slightly, showing that continual update of the learning parameters (as the system equilibrates) is possible as long as the learning rate is low enough [15, 16, 17].

Certain natural neural networks, however, do not require that the learning rate is slow compared to equilibration of physical degrees of freedom. They reach a steady spiking state on a similar timescale

[*]nachis@sas.upenn.edu

Fourth Workshop on Machine Learning and the Physical Sciences (NeurIPS 2021).

with synaptic plasticity processes [18, 19, 20]. If the environment supplying the training signals changes frequently, and the learning system must quickly adapt to the changing environment, it is sensible for learning to occur on similar timescales to that of the system response (equilibration of the physical degrees of freedom). This may also be the case in the humoral immune response, where antibodies are formed rapidly after infection, even before the onset of symptoms [21, 22].

Here we show that learning is highly effective even when equilibration rates of the physical and learning degrees of freedom are comparable. We find that for simple allosteric learning tasks [23] (where target nodes produce desired outputs when inputs are applied at source nodes) in linear resistor networks, learning can be made *faster* than the physical dynamics. This has the effect of speeding up network training in real time by orders of magnitude. We show experimental and numerical evidence for a critical learning rate beyond which learning dynamics become oscillatory, yet still successful.

2 Out of equilibrium coupled learning

In most contrastive learning approaches, physical degrees of freedom x_a first equilibrate, and then the equilibrium states (denoted by stars) x_a^* are used to modify learning degrees of freedom w_j. Coupled learning [12] compares two equilibrium states (free state x_a^{*F} and clamped state x_a^{*C}) to derive local learning rules (approximating the gradient of the cost function). Suppose that a network trains on an allosteric task, in which when a set of constraints are applied to source nodes X_S, the network responds by having desired values at (typically physically distant) target nodes X_T. Such allosteric effects are not generic in physical networks, but are common in biology, for example in proteins (mechanical allostery) and vasculature (flow allostery) [23]. A free state x_a^{*F} is defined by applying source constraints X_S and letting the network equilibrate, reading off the target values x_T^{*F}. Then, a supervisor clamps the target nodes, nudging them toward the desired output $x_T^C = x_T^{*F} + \eta[X_T - x_T^{*F}]$, with $\eta \ll 1$. The coupled learning rule is given by

$$\dot{w}_j = \alpha\eta^{-1}\frac{\partial}{\partial w_j}\{E^{*F}(X_S, x_T^{*F}) - E^{*C}(X_S, x_T^C)\}, \tag{1}$$

where E^{*F} and E^{*C} are the energies (or dissipated power) of the equilibrium free and clamped states, and α a learning rate. The contrastive function $\mathcal{C}^* \equiv \eta^{-1}[E^{*C} - E^{*F}]$ was shown to approximate the mean squared error cost function $C^* \sim (x_T^{*F} - X_T)^2$ in the limit $\eta \to 0$ [12].

The coupled learning rule above assumes both the free and clamped states are in equilibrium (or steady state). Suppose instead that the network learns at the same time as it attempts to equilibrate, following its physical dynamics, e.g. overdamped dynamics:

$$\dot{\boldsymbol{x}}^F = -\gamma\partial_{\boldsymbol{x}}E^F \quad ; \quad \dot{\boldsymbol{w}} = -\alpha\partial_{\boldsymbol{w}}\mathcal{C}, \tag{2}$$

α/γ sets the ratio between learning and physical rates ($\alpha/\gamma \to 0$ is the quasi-static limit of coupled learning). Note that computing Eq. 2 is linear in network size, so the computational efficiency of learning is comparable to back-propagation on feed-forward neural networks.

We study the special case of linear electric resistor networks. Here, the physical degrees of freedom are voltages at nodes v_a, and the learning degrees of freedom are edge conductances (inverse resistances) $g_j = R_j^{-1}$. The voltage drop over each edge is the difference in voltage between the two nodes connected by that edge $\Delta v_j \equiv \Delta_{ja} v_a$, with Δ_{ja} the incidence matrix. The power dissipation is given by a sum over edges $\mathcal{P} = \frac{1}{2}\sum_j g_j \Delta v_j^2$. The network steady state is given by voltages v_a minimizing the power dissipation $\mathcal{P}$, subject to constraints. In ideal networks without capacitance, the network reaches its steady state instantaneously. Once capacitance is relevant, it decreases the physical relaxation rate γ, and we can simulate the system, e.g. by postulating overdamped dynamics

$$\begin{aligned} \dot{v}_a^{F,C} &= -\gamma \sum_j \Delta_{aj}^T g_j \Delta_{ja} v_a^{F,C} \\ \dot{g}_j &= \alpha\eta^{-1}\{[\Delta v_j^F]^2 - [\Delta v_j^C]^2\}, \end{aligned} \tag{3}$$

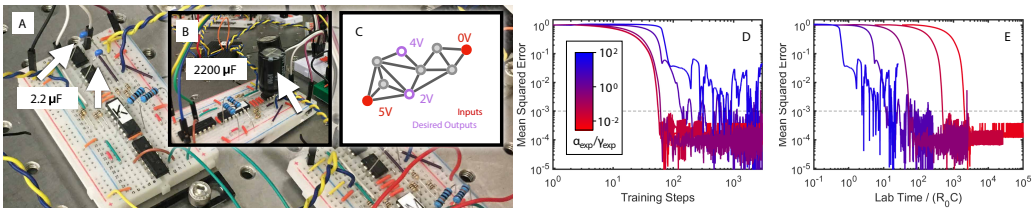

Figure 1: Out of equilibrium learning in a physical network. (A) A single edge (in both free and clamped networks) with parallel 2.2μF capacitors highlighted with arrows. (B) A single edge with a parallel 2200μF capacitors. (C) Network structure with inputs and outputs for the allosteric task used in (D)-(E). Typical error traces for an allosteric task as a function of training steps (D) and relaxation times (E). Colors indicate ratio of learning to relaxation rates.

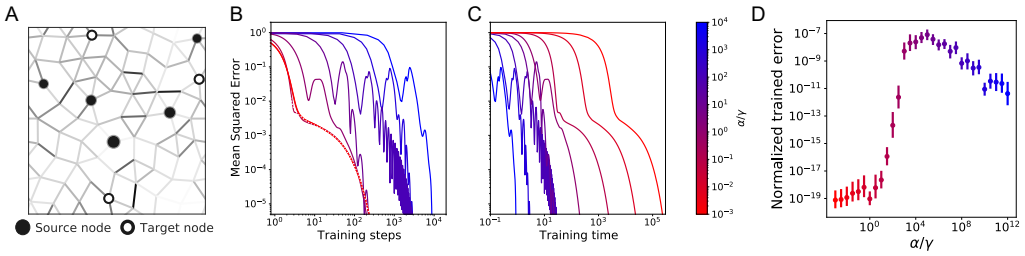

Figure 2: Out of equilibrium learning in simulated resistor networks. (A) Disordered resistor network (64 nodes, 5, sources, 3, targets). Full (empty) circles denote source (target) nodes, dot sizes showing the magnitudes of the source voltages V_S (target voltages V_T). Edge thickness indicates power through each resistor. (B) Error during training at different ratios of learning to physical rates α/γ. When $\alpha \gtrsim \gamma$, more training steps are required to reduce the error (dotted line - $\alpha/\gamma \to 0$). (C) Higher learning rate results in more rapid learning in real time. (D) The error in trained networks does not change up to $\alpha \approx \gamma$. Even beyond $\alpha \approx \gamma$, learning greatly reduces errors far from equilibrium.

Physical experiments.– To study if physical and learning dynamics can be coupled in the way suggested by Eqs. 2-3, we perform experimental and numerical studies. In recent work, we demonstrated a variable electric resistor learning network [24] based on coupled learning [12]. The experiment operated in the quasi-static regime, where learning was far slower than the network response times. To relax the quasi-static condition, we slow down the physical relaxation within such a network using capacitors in parallel to each variable resistor, as shown in Fig. 1AB. The charging of capacitors increases the time to reach steady current in the network. During the experiment a pause $\tau/2$ prior to each learning step allows the system to partially equilibrate before updating its resistance value. The learning rate is $\alpha_{exp} = \tau^{-1}$ and the physical relaxation rate is $\gamma_{exp} \approx (R_0 c)^{-1}$, where c is the capacitance parallel to each resistor (2.2, 22, 220, or 2200 μF). We estimate the relaxation rate using resistance $R_0 \equiv 10K\Omega$ (each resistor varies between $781\Omega - 100K\Omega$).

The system is trained for a 2-target, 2-source allosteric task (Fig. 1C) as capacitance and τ are varied. We measure the instantaneous mean-squared error, i.e. the squared difference between the free state outputs and the desired outputs, normalized by its initial value. When physical relaxation is slow compared to the learning rate ($\alpha_{exp}/\gamma_{exp} \ll 1$) the system learns the task and the error hovers at the noise floor (Fig. 1DE). In this regime α/γ does not affect the training time when measured in training steps (Fig. 1D) but does when measured in time (Fig. 1E). As the learning rate is increased, the system's behavior changes markedly, at first taking more training steps to learn, and then entering into a regime with perpetual oscillations ($\alpha_{exp}/\gamma_{exp} \gtrsim 1$), as seen best by the blue curve in Fig. 1.

Numerical experiments.– In simulations, we train a disordered resistor network (64 nodes, 143 edges, code available [25]) for an allosteric task involving 5 source nodes with voltages sampled from $V_S \sim \mathcal{N}(0,1)$, and 3 target nodes with desired voltages $V_T \sim 0.2\mathcal{N}(0,1)$ (Fig. 2A). During training using Eq. 3 ($\eta = 10^{-3}$), we measure the instantaneous mean squared error $C \equiv \sum_T (V_T - v_T^F)^2$ (normalized at its initial value). In Fig. 2B, error values are shown for different ratios α/γ as a function of training steps. At slow learning $\alpha/\gamma \to 0$ we recover the quasi-static case of coupled learning (white dotted line). The learning rate α can be increased substantially to $\alpha \approx \gamma$ with little

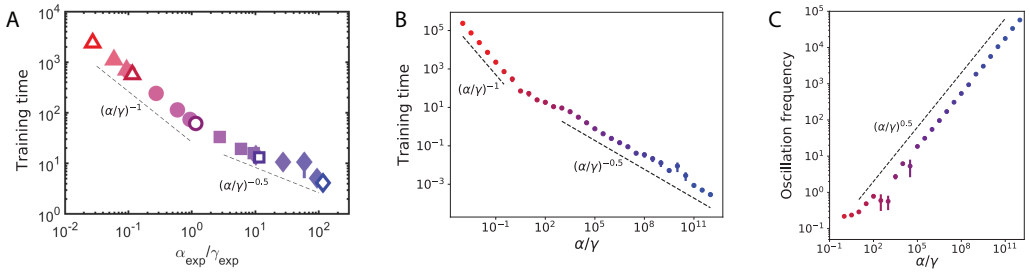

Figure 3: Training times, i.e. time for normalized error to reach a threshold $\tilde{C} = 10^{-3}$ in (A) experiments and (B) simulations. A transition from near equilibrium dynamics to strongly out of equilibrium dynamics occurs at $\alpha_c \sim \gamma$. For $\alpha < \gamma$, training time improves linearly with α. When $\alpha \gg \gamma$ networks are trained far from equilibrium, with training times improving as $\sqrt{\alpha}$. (C) Starting at the critical learning rate $\alpha_c \approx \gamma$, error exhibits oscillations whose frequency grows as $\sqrt{\alpha/\gamma}$.

effect on learning dynamics. Further increasing α, the error oscillates at short times, yet the network still learns successfully. Notably, simulations recover the experimental phenomenology despite multiple differences in how these systems learn, suggesting that Eq. 2 captures essential aspects of physical learning out of equilibrium.

In physical systems, where γ is fixed by the network physics, slow learning rates entail long training times. In Fig. 2C we set the time scale by $\gamma = 1$, finding that increasing the learning rate, even far beyond the physical rate, significantly reduces training times. In Fig. 2D, we plot the error for trained networks of size 64, with 10 sources and 3 targets, averaged over 50 tasks. We see that the error of these networks after training is largely unaffected for $\alpha < \gamma$. At higher α, errors are larger but still small; networks can be trained effectively for allosteric tasks far beyond equilibrium. This transition to strong out of equilibrium dynamics is marked by the onset of error oscillations (Fig. 2BC).

Out of equilibrium effects also manifest through a change in scaling of the training time. In Fig. 3AB, we plot the time taken the networks to reach a normalized error threshold $\tilde{C}(T) = 10^{-3}$. For quasi-static training $\alpha \ll \gamma$ we see a linear improvement in the training time $T \sim (\alpha/\gamma)^{-1}$. Transitioning to out of equilibrium learning, training becomes faster, but with slower scaling $(\alpha/\gamma)^{-\frac{1}{2}}$.

Origin of learning oscillations.– How do networks learn even far from equilibrium? What gives rise to learning oscillations at $\alpha \gtrsim \gamma$? We can gain insight on these questions by taking a second time derivative of the physical degrees of freedom in Eq. 2:

$$\gamma^{-2}\delta\ddot{x}^F + \gamma^{-1}H\delta\dot{x}^F + \frac{\alpha}{\gamma}D^2\delta x^F = 0, \tag{4}$$

where we have defined the difference between the observed and desired states $\delta x^F \equiv x^F - X$, and the physical Hessian and cross derivatives $H \equiv \partial_x^2 E^F, D \equiv \partial_x \partial_w E^F$. These dynamics can be partitioned to an overdamped term for physical relaxation and an oscillatory term which scales with the learning rate. As H, D do not depend on α, γ, the oscillatory term becomes important only when learning is fast enough. The solutions to this equation become oscillatory when the discriminant $H^2 - 4\frac{\alpha}{\gamma}D^2$ becomes negative at $\alpha_c \sim \gamma$. Beyond this critical learning rate the oscillation frequency increases, scaling as $\sqrt{\alpha/\gamma}$. This frequency scaling then dictates the training time, which can be estimated by the first pass of the oscillatory dynamics through zero. These predictions are verified in Fig. 3C for 64 networks with 10 sources and 3 targets. The oscillatory learning dynamics converge to solutions $\delta x \to 0$ as long as the physical derivatives H, D vary slowly relative to the physical and learning degrees of freedom x, w, i.e. if the energy landscape is smooth enough.

Eq. 4 also explains the observed power laws in the training time. At low learning rates error convergence is dominated by the slowest overdamped decay mode of H, scaling as $\tau \sim \gamma/\alpha$. Far from equilibrium however, where $\alpha \gg \gamma$, convergence is dominated by the inverse oscillation frequency, scaling as $\tau \sim \sqrt{\gamma/\alpha}$.

3 Concluding remarks

Going beyond quasi-static contrastive learning, we observe that physical networks can learn out-of-equilibrium. This idea may be important in biological learning, where plasticity time scales of some neural circuits are similar to that of reaching steady state [18, 20]. In machine learning, equilibrating recurrent networks constitutes a computational bottleneck [14]. As full equilibration is not required, the computational load of training recurrent networks may be significantly reduced.

Our results show that trained physical networks exhibit overdamped learning dynamics until a critical learning rate similar to the physical rate $\alpha_c \sim \gamma$. Increasing the learning rate to α_c can speed up learning by orders of magnitude compared to quasi-static training. Beyond the critical threshold, we describe the formation of under-damped oscillations in the learning dynamics, resulting from the learning parameters varying too quickly and overshooting desired solutions. In complex non-convex landscapes underlying some learning problems, one would expect training far from equilibrium to fail as the network rapidly jumps between energy basins due to the changing learning parameters. We show that at least in the linear flow networks we trained for allosteric tasks, in which the learning landscape is convex, training is possible and successful far from equilibrium.

Broader Impact

Local physically plausible learning rules based on contrastive learning have emerged as possible candidates for constructing physical "learning machines", able to autonomously adapt to external influence in order to gain desired functionality. However, such rules typically rely on equilibration of the physical state before any learning update, so such systems are prone to very slow training. We break the assumption of quasi-static learning in physical systems, showing that they can be trained at comparable rates to their physical relaxation, similar to certain biological neural circuits. This property is crucial for training of slow physical systems, as well as using such contrastive learning methods as machine learning algorithms for recurrent networks.

Acknowledgments and Disclosure of Funding

This research was supported by the U.S. Department of Energy, Office of Basic Energy Sciences, Division of Materials Sciences and Engineering award DE-SC0020963 (M.S.), the Army Research Office award W911NF-21-1-0076 and air force office of scientific research award (FA9550-21-1-0313) (M.Z.M.), and the National Science Foundation via the UPenn MRSEC/DMR-1720530 (S.D. and D.J.D.), and DMR-2005749 (A.J.L.).

References

[1] Kyungdoc Kim, Seokho Kang, Jiho Yoo, Youngchun Kwon, Youngmin Nam, Dongseon Lee, Inkoo Kim, Youn-Suk Choi, Yongsik Jung, Sangmo Kim, et al. Deep-learning-based inverse design model for intelligent discovery of organic molecules. *npj Computational Materials*, 4(1):1–7, 2018.

[2] Miguel A Bessa, Piotr Glowacki, and Michael Houlder. Bayesian machine learning in metamaterial design: Fragile becomes supercompressible. *Advanced Materials*, 31(48):1904845, 2019.

[3] Juhwan Noh, Jaehoon Kim, Helge S Stein, Benjamin Sanchez-Lengeling, John M Gregoire, Alan Aspuru-Guzik, and Yousung Jung. Inverse design of solid-state materials via a continuous representation. *Matter*, 1(5):1370–1384, 2019.

[4] Wei Ma, Feng Cheng, Yihao Xu, Qinlong Wen, and Yongmin Liu. Probabilistic representation and inverse design of metamaterials based on a deep generative model with semi-supervised learning strategy. *Advanced Materials*, 31(35):1901111, 2019.

[5] Paul Z Hanakata, Ekin D Cubuk, David K Campbell, and Harold S Park. Forward and inverse design of kirigami via supervised autoencoder. *Physical Review Research*, 2(4):042006, 2020.

[6] Nidhi Pashine, Daniel Hexner, Andrea J Liu, and Sidney R Nagel. Directed aging, memory, and nature's greed. *Science Advances*, 5(12):eaax4215, 2019.

[7] Daniel Hexner, Nidhi Pashine, Andrea J Liu, and Sidney R Nagel. Effect of directed aging on nonlinear elasticity and memory formation in a material. *Physical Review Research*, 2(4):043231, 2020.

[8] Menachem Stern, Chukwunonso Arinze, Leron Perez, Stephanie E. Palmer, and Arvind Murugan. Supervised learning through physical changes in a mechanical system. *Proceedings of the National Academy of Sciences*, 117(26):14843–14850, 2020.

[9] Javier R Movellan. Contrastive hebbian learning in the continuous hopfield model. In *Connectionist models*, pages 10–17. Elsevier, 1991.

[10] Benjamin Scellier and Yoshua Bengio. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience*, 11:24, 2017.

[11] Jack Kendall, Ross Pantone, Kalpana Manickavasagam, Yoshua Bengio, and Benjamin Scellier. Training end-to-end analog neural networks with equilibrium propagation. *arXiv preprint arXiv:2006.01981*, 2020.

[12] Menachem Stern, Daniel Hexner, Jason W Rocks, and Andrea J Liu. Supervised learning in physical networks: From machine learning to learning machines. *Physical Review X*, 11(2):021045, 2021.

[13] Victor Lopez-Pastor and Florian Marquardt. Self-learning machines based on hamiltonian echo backpropagation. *arXiv preprint arXiv:2103.04992*, 2021.

[14] Sergey Bartunov, Adam Santoro, Blake Richards, Luke Marris, Geoffrey E Hinton, and Timothy Lillicrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In *Advances in Neural Information Processing Systems*, pages 9368–9378, 2018.

[15] Maxence Ernoult, Julie Grollier, Damien Querlioz, Yoshua Bengio, and Benjamin Scellier. Updates of equilibrium prop match gradients of backprop through time in an rnn with static input. *arXiv preprint arXiv:1905.13633*, 2019.

[16] Maxence Ernoult, Julie Grollier, Damien Querlioz, Yoshua Bengio, and Benjamin Scellier. Equilibrium propagation with continual weight updates. *arXiv preprint arXiv:2005.04168*, 2020.

[17] Benjamin Scellier. A deep learning theory for neural networks grounded in physics. *arXiv preprint arXiv:2103.09985*, 2021.

[18] Robert S Zucker and Wade G Regehr. Short-term synaptic plasticity. *Annual review of physiology*, 64(1):355–405, 2002.

[19] Ryota Kanai and Frans AJ Verstraten. Perceptual manifestations of fast neural plasticity: Motion priming, rapid motion aftereffect and perceptual sensitization. *Vision research*, 45(25-26):3109–3116, 2005.

[20] Yaniv Sagi, Ido Tavor, Shir Hofstetter, Shimrit Tzur-Moryosef, Tamar Blumenfeld-Katzir, and Yaniv Assaf. Learning in the fast lane: new insights into neuroplasticity. *Neuron*, 73(6):1195–1203, 2012.

[21] World Health Organization et al. *WHO guidelines on hepatitis B and C testing*. World Health Organization, 2017.

[22] Jangu E Banatvala and David WG Brown. Rubella. *The Lancet*, 363(9415):1127–1137, 2004.

[23] Jason W Rocks. *Allosteric Functionality In Mechanical And Flow Networks*. PhD thesis, University of Pennsylvania, 2019.

[24] Sam Dillavou, Menachem Stern, Andrea J Liu, and Douglas J Durian. Demonstration of decentralized, physics-driven learning. *arXiv preprint arXiv:2108.00275*, 2021.

[25] Menachem Stern. Source code: Out of equilibrium coupled learning. https://github.com/NachiStern/OutOfEq-Coupled-Learning, 2021.

Paper checklist

Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**

Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**

Did you discuss any potential negative societal impacts of your work? **We discuss speeding up learning in physical networks. We do not see a specific way in which this work can have negative impact beyond other physically plausible learning algorithms.**

Did you describe the limitations of your work? **Yes. We study out of equilibrium dynamics in linear resistor networks learning simple allosteric tasks. We argue on how training may fail in more complicated settings.**

Did you state the full set of assumptions of all theoretical results? **Yes**

Did you include complete proofs of all theoretical results? **No. Our work is not concerned with mathematical proofs, just physical experiments and computer simulations.**

Did you include the code, data, and instructions needed to reproduce the main experimental results? **Yes**

Did you specify all the training details? **Yes**

Did you report error bars? **Yes**

Did you include the amount of compute and the type of resources used? **Minor computation resources used on internal cluster. Experimental resources include one electric resistor system with a power consumption significantly lower than a desktop computer.**