Simulated-annealing-based adaptive equaliser for on-die variation compensation

H. Chung and G.-Y. Wei

Fully exploiting the flexibility of lookup-table-based equalisers, it is proposed to compensate for on-die variation effects within a transmit-side equaliser. To efficiently deal with the nonlinear nature of circuit non-idealities, the proposed equaliser utilises simulated annealing for adaptation.

Introduction: As processes continue to scale, worsening on-die parameter variation exacerbates circuit/system performance and reliability [1, 2]. For the case of high-speed backplane transmitters, device mismatch adds offsets to analogue voltages and signal phases, leading to signal integrity degradation [3]. It is important to minimise such offsets, especially at high data rates, as transmitter jitter gets amplified by channel bandwidth limitations [4, 5]. Lookup-table-based (LUT-based) transmit-side equalisers have been proposed [6, 7] to accommodate wide ranges of signalling schemes and channel environments, utilising LUTs' flexibility to reprogram equalisation settings and capability to support both linear and nonlinear equalisations. Fortunately, such advantages of LUTs also enable an LUT-based equaliser to efficiently deal with on-die variation effects that are unpredictable prior to fabrication and nonlinear. In [3], it has been shown that voltage and phase offsets due to on-die variations can be addressed within a LUT-based equaliser. However, the equalisation scheme in [3] assumes that all the circuit non-idealities are known in advance, which is impractical. In this Letter, we propose an adaptive equalisation scheme that utilises simulated annealing (SA) for circuit and channel non-idealities compensation.

SA-based equalisation: To realise an adaptive equalisation, the LUT-based equaliser requires an algorithm that can efficiently search a given solution space (i.e. possible LUT settings). Unfortunately, the solution space can be large, especially when the number of taps and bit resolutions are high, as the number of possible LUT settings grows exponentially with the number taps and bit resolutions. Furthermore, nonlinear offsets due to on-die variations roughen the large solution space, which may hinder convergence of the equaliser. Therefore an equalisation algorithm that can efficiently deal with such a large and rough solution space is needed.

SA [8] is an optimisation algorithm that works well for systems with a large and rough solution space. An adaptive equalisation based on SA iterates through a solution space in a probabilistic manner to minimise an energy function that represents the desired attributes of the transmitted signal. As it operates in a probabilistic manner, the SA-based equaliser does not always guarantee a single optimal equalisation setting. However, one single solution may not exist for the rough solution space worsened by on-die variations. Rather, the annealing process constantly iterates through the space to find close-to-optimal equalisation settings, which can also change over time owing to slow environmental parameter variations.

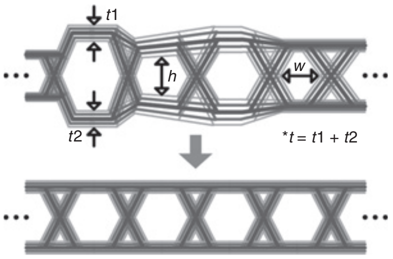

Fig. 1 *Illustration of energy function parameters*

To define an energy function, we assume a time-interleaved transmitter structure, presented in [3]. Fig. 1 shows an example of time-interleaved eyes at the transmitter output and their attributes used in the energy function. Phase and voltage offsets due to on-die variation result in uneven eye openings across the interleaved eyes [3], significantly degrading the transmitter performance as the overall performance depends on the worst-case eye opening. Therefore, the energy function focuses on evening out the interleaved eyes. First, since the base clock period is a bounded sum of all the interleaved eye widths, an adaptation scheme that always maximises the worst-case eye width eventually leads to evenly distributed widths and minimises jitter. Eye height, on the other hand, is not bounded and adds two components to the energy function: one to maximise the worst-case height and the other to even out the heights (i.e. flatness) across all interleaved eyes. Lastly, the energy function includes a component to reduce variation in high and low voltage levels (i.e. thickness) of each symbol, which helps to reduce residual intersymbol interference (ISI). The resulting energy function is:

$$\text{energy} = 1 - \frac{\min(w)}{a1} - \frac{\min(h)}{a2} + \frac{\text{flat}(h)}{a3} + \frac{\max(t)}{a4}.$$

where constants $a1-a4$ normalise the components and scale them according to relative importance.

The iteration process works by (*a*) adding small perturbations to the equaliser settings, (*b*) capturing the transmitted signal, (*c*) computing the energy, (*d*) comparing the computed energy to an optimum energy and (*e*) updating the optimum energy and equaliser settings if the computed energy is lower. The algorithm can escape local minima by accepting non-optimal settings with a probability inversely proportional to a system temperature, which represents the annealing process. In other words, the algorithm is more likely to accept non-optimal settings in the beginning (i.e. when the system temperature is high), but this probability reduces as the system anneals over time. The probability function and the condition when nonoptimal equalisation settings are accepted are:

$$\text{probability} = \exp\left(\frac{\alpha \times (energy_{curr} - energy_{opt})}{T}\right) \geq r$$

where α is a scaling factor, $energy_{curr}$ is current energy, $energy_{opt}$ is the optimal energy at the moment, T is the system temperature and r is a random number between 0 and 1. The scaling factor α is selected such that roughly 50% of non-optimal settings are accepted when T is high and almost none (i.e. 0.01%) is taken when T is low.

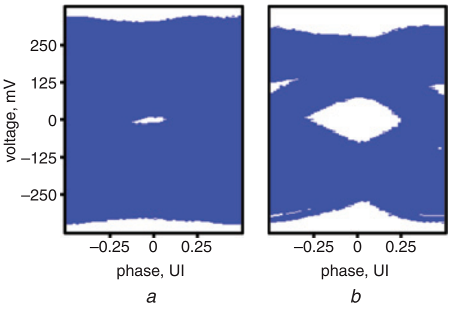

Fig. 2 *Measured eye diagrams at 11.2 Gbit/s*

a Before equalisation
b After equalisation

Experimental results: The proposed SA-based equalisation has been verified with an eight-way time-interleaved transmitter test-chip prototype fabricated in a 0.13 μm CMOS process [3]. An external sampling oscilloscope was employed to monitor transmitted eye attributes used for the proposed SA-based equalisation.

Fig. 2 shows the transmitted eye diagrams before (Fig. 2*a*) and after (Fig. 2*b*) the proposed SA-based equalisation at 11.2 Gbit/s. Eight time-interleaved eyes are overlaid to show a final aggregate eye opening. In this experiment, the equalisation process mostly focuses on maximising the worst-case eye width (i.e. $a1 = 0.5$) to even out interleaved eye widths and other eye attributes have relatively low importance (i.e. $a2 - a4 = 0.167$). Fig. 2*b* demonstrates that the proposed SA-based equalisation successively opens up the almost closed eye by compensating for phase and voltage offsets due to on-die variations.

Fig. 3 quantitatively compares eye attributes across the eight time-interleaved eyes and confirms that SA-based equalisation can successfully even out eye widths and heights, maximising the worst-case eye opening. Experimental results with the oversampled zero-forcing equalisation [3], which can be considered as an optimal solution, suggest that the proposed SA-based equalisation can find close-to-optimal equaliser settings through the adaptation process.