


GRAPHICS PROCESSING UNITS

## Just about managing

*IEEE Trans. Comput.* [https://doi.org/10.1109/](https://doi.org/10.1109/TC.2018.2789904)

TC.2018.2789904 (2018)



Credit: IEEE

Graphics processing units (GPUs) are used to handle computationally intensive tasks such as 3D graphics rendering or cryptocurrency mining. In contrast to central processing units (CPUs) found in household computers, a GPU consists of many cores, typically on the order of hundreds to thousands, which allow massively parallel processing of multiple software instructions. Within each core, however, process variations due to small feature sizes, chip aging and within-die parameter deviations can exist. Over time, these cores perform differently under stress and, without proper management, may contribute to a notable drop in chip-level performance.

Haeseung Lee and colleagues at the University of California Irvine and TU Wien have now developed an approach to

manage the resources of embedded GPUs that takes into account the current aging status and process variations of each core in order to minimize GPU aging. The key is to first understand the stress distribution of the cores and then balance the workload across the cores by assigning a different number of instructions to the clusters. Their approach improves GPU aging in over 95% of cases (state-of-the-art compiler-based techniques improve GPU aging in 72.25% of cases). Moreover, compared to state-of-the-art compiler-based techniques, their approach improves performance overhead by 40%.

---

Michael Lee

Published online: 8 February 2018  
<https://doi.org/10.1038/s41928-018-0035-y>