Energy Aware Algorithm Design via Probabilistic Computing: From Algorithms and Models to Moore's Law and Novel (Semiconductor) Devices[1]

Krishna V. Palem
Georgia Institute of Technology

The energy consumed by computations is a significant concern, especially within the context of embedded systems, on par with the past focus on raw speed or its derivative, performance in the high-performance computing domain. In this talk, we will outline an entirely new approach to energy-aware computing: trading the probability of the BIT being correct for savings in the energy consumed, yielding a *probabilistic* bit or PBIT (instead of a conventional BIT which is guaranteed to be correct). At its heart, the approach taken here is built on the fundamental and novel thesis that the energy consumed by a computation is proportional to the associated *accuracy,* characterized as the probability of being correct, with which each ``bit'' is computed. With this as background, *probabilistic* hardware devices---these can be viewed as the ``hardware'' counterparts of the well-known probabilistic algorithms---and gates realized from conventional CMOS technology for computing PBITs will be described. Our probabilistic devices are constructed through the counterintuitive approach of using *noise,* which is increasingly being viewed as a hurdle to sustaining Moore's law, as a resource rather than as an impediment. Specifically, we have demonstrated that coupling thermally induced sources of noise, as well as the prevalent *power-supply* noise with a conventional CMOS device yields a *probabilistic* switch, which can in turn be a basis for realizing probabilistic applications in silicon. These probabilistic (hardware) switches compute with a definite probability of error, and have been demonstrated to serve as natural building-blocks in architectures for supporting probabilistic algorithms, yielding significant savings to the (*energy* x *performance*) metric in a variety of embedded computing applications ranging over speech and pattern recognition, robotics and others---improvements of over a factor of 100 within the context of an AMI 0.5μm, a TSMC 0.25μm and proprietary deep submicron processes, when compared to executing the same applications on a low energy embedded processor, the StrongARM SA-1100. At a deeper level, all of this work rests on the twin foundations of classical thermodynamics (of Maxwell, Boltzmann and Gibbs), and the relatively modern computational complexity theory. Time permitting, these foundations will be surveyed.

Krishna V. Palem holds Professorships with tenure in Electrical and Computer Engineering and in Computer Science in the College of Computing, a senior research leadership in the College of Engineering, and has been the founding director of the *Center for Research in Embedded Systems and Technology* (CREST) (www.crest.gatech.edu) at the Georgia Institute of Technology, since 1999. Previously, he held positions at the IBM T. J. Watson Research Center and NYU's Courant Institute of Mathematical Sciences (Computer Science). His work is recognized internationally in academia and in industry ranging from contributions to algorithms, compiler optimizations, reconfigurable computing systems, as well as in power-aware computing and most recently microelectronics. Over the past decade, he has focused on applying innovations from these disciplines to the increasingly significant domain of embedded computing systems. He has awards for excellence from Hewlett Packard, IBM and Panasonic and among others serves on the editorial boards of the recently formed *ACM Transactions on Embedded Computing Systems*. Palem laid the foundations of *architecture assembly* which is at the heart of the product offerings of Proceler Inc.---an Atlanta based venture. The prestigious Analysts' Choice Awards recognized Proceler's technology, by nominating it as one of the outstanding technologies of 2002*,* and his Ph.D advise Suren Talla was recognized with a dissertation award from NYU for aspects of this work. He has chaired (and co-chaired) meetings whose advise has led to funding initiatives in Embedded and Hybrid Systems in the US, as well as by the leading research funding agency A*Star of Singapore. He was a Schonbrunn visiting professor at the Hebrew University of Jerusalem, Israel, where he was recognized for *excellence in teaching*. . He is a Fellow of the IEEE.

---

[1] This work was supported in part by DARPA Seedling Contract #F30602-02-2-0124.