1.1 Trends in Power Consumption for Computation

Humans have wanted an *internet of things* (IoT) smart, battery-less objects exchanging information with each other and with people for a very long time in popular science articles and science fiction stories. Only now, however, do we have the technology to realize power-stingy operation, energy-harvesting, and incredibly low-powered exchanges of information. Trends in electronics indicate that we will soon start building this science fiction.

In the world of digital electronics, perhaps the most striking trend over the last 70 years has been the inexhorable march towards lower and lower power consumption. Koomey published an informative summary of these trends specifically for computer platforms, demonstrating how the number of computational operations per unit of energy consumed has been exponentially decreasing for 70 years [**?**]. This trend is illustrated in the left-hand graph of Figure 1.1, illustrating the average decline as a function of year over this time period. The result is particularly surprising when once considers how constant these reductions in computational power consumption have been achieved over different epochs of computing hardware: the trendline is nearly constant whether vacuum tubes, discrete transistors, microprocessor integrated circuits, or multi-core processors are being used for the computation.

A large portion of the microprocessor-era gains in energy efficiency were achieved by reducing the size of individual transistors on integrated circuits. Smaller transistors required less current and voltage to change states, thereby allowing significant computational energy reductions. The smaller devices could also be operated *faster*, resulting in *Moore's Law* (an observation by Intel's co-founder, Gordon Moore) which stated that processor speeds seemed to double every two years. This "law" eventually exhausted in the early 2000s due to thermal limitations of integrated circuit devices: the volume of devices shrunk faster than the power consumption, meaning that tinier devices ultimately struggled with heat dissipation.

Koomey's Trend in Figure 1.1 is a far more generalized version of Moore's Law. Clearly gains in the energy-efficiency of computation before and after the basic integrated circuit microprocessor era have been achieved, implying that size reduction of transistors on integrated circuits is only partly responsible. Furthermore, there appears to be enough innovation in the research pipeline to continue reduction of computational power consumption to 2020 and beyond [**?**]. Particular innovations in computer architectures, subthreshold logic devices, asynchronous logic, and even quantum computing could push computational efficiency much further.

In many ways, the microprocessor industry drives much of the innovation in electronics. The huge microprocessor commercial market – valued over 92 billion USD as of 2012 (http://www.itcandor.com/microprocessor-q312/) and growing – invests significantly into improved chip manufacturing. Other lower-end digital electronic devices – microcontrollers, RFID tags, microelectromechanical sytem (MEMS) sensors, and small displays, to name just a few examples – must partly rely on the innovation of the microprocessor industry for improvements in energy efficiency. As

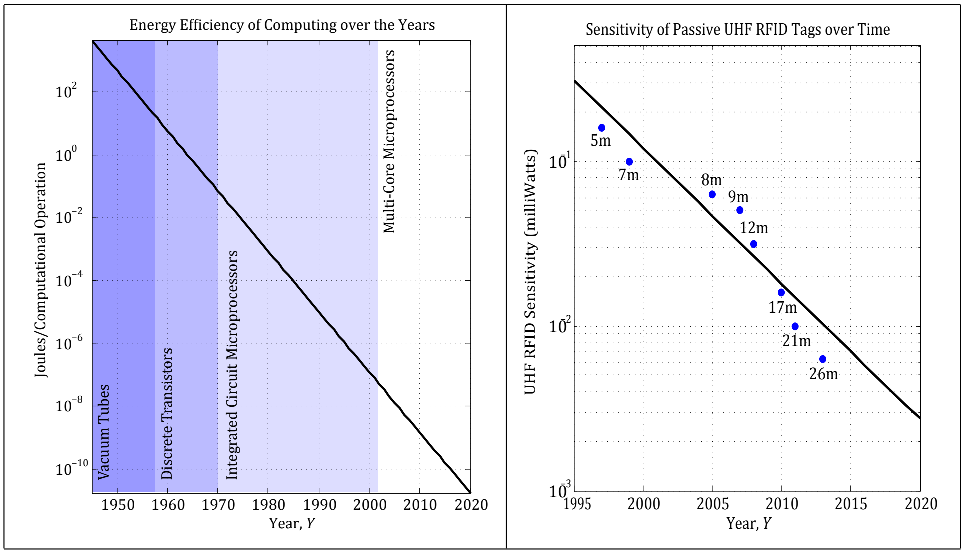

Figure 1.1. Trends from two different industries demonstrating exponential reductions in computational efficiency. Koomey's trend for improved computational energy efficiency for computer platforms is illustrated on the left [?], while Nikitin's trend for improved sensitivity and range of passive UHF RFID tags is illustrated on the right [?].

a result, there is often a time lag between when innovation in the microprocessor world makes its way into smaller-market electronic device. For example, microprocessors implementing transister device sizes of 180 nm were sold commercially in 1999; the first commercial UHF RFID tag using 180nm logic was sold commercially in 2010.

Passive UHF RFID is another example of an industry where energy-efficiency gains in computation and communication are exponential. In passive RFID, an RFID *tag* that has no on-board power source communicates back to a *reader* that is illuminating the tag with RF power – but only if a minimum amount of power is received. This minimum power required is the RFID tag's *sensitivity.* On the right-hand side of Figure 1.1 is an illustration of how the sensitivity of passive RFID tags has decreased since the mid-1990s, with selective data points adapted from Nikitin's observations in [?]. As the sensitivity drops, a passive RFID tag can operate further and further from the reader for fixed illumination conditions. Estimates of the tag-reader separation distance using a typical reader operating in an ideal radio channel are included next to the selected data points of Figure 1.1.

Nikitin's Trend in Figure 1.1 again illustrates the potential for further reductions in sensitivity and energy-efficiency of passive RFID tags. Since these types of

RFID tags represent the cutting-edge radio devices in terms of energy consumption per bit of transmitted data, their study will be an important part of this course. Prognostication of the sensitivity, S, for these types of devices can be estimated as a function of year, Y, by applying this simplified empirical expression:

$$S = \left(\frac{6}{5}\right)^{1989-Y} \text{ milliWatts} \tag{1.1.1}$$

Compare Equation (1.1.1) to a simplified empirical expression of Koomey's Trend:

$$E = \left(\frac{14}{9}\right)^{1964-Y} \text{ Joules/operation} \tag{1.1.2}$$

Both trends indicate a future where extraordinarily low-powered devices can perform basic sensing, computation, and radio communication over distance using extraordinarily limited amounts of energy.