Fighting for a Niche: An Evolutionary Model of Storage

Avani Wildani
Emory University

We know intuitively that if two devices, such as HDVD and Blu-Ray discs, have no clear difference in terms of cost, speed, or capacity, one will eventually leave the market. In evolutionary biology, this principle is termed "competitive exclusion": if species occupy the same environmental niche, one will eventually outcompete the others (though predicting which will live is out of scope, we can see that the niche can only support one). Our wild and crazy plan is to project the future of storage systems by tracking the co-evolution of devices along with the underlying storage marketplace, or our "niche."

While technical challenges such as bit density are a factor, it is infamously hard to predict when new technologies will take off in the storage market. As a community, we've been talking about phase change memory, MRAM, racetrack memory, memristors, and other "upcoming" storage devices for a generation. While there have been significant theoretical and engineering advances in many of these device types, they have for the most part not hit the mass market. Discussions of why this happens tend to devolve into either focusing on a minor technical hurdle or blaming some combination of "industry foresight," "market pressure," or the ever nebulous "pipeline."

Any project that attempts to address archival storage must make some assumptions about the future device landscape. These models are often forced to make naïve assumptions about the trajectory of storage devices using simple linear scaling or Kryder's rate (which assumes disks will improve in density by 15% per year), and all device types are presumed to be independent [2]. Moreover, projections typically only consider a single device attribute, such as cost per byte, even though devices clearly grow in different dimensions at different rates over time. The classic example of this is the lack of improvement in disk speed over a time period where cost has improved dramatically. Figure 1 shows a set of devices including disks, SSDs, and flash sticks, where the color represents the release date. Note that the devices, instead of clustering across time, cluster by time.

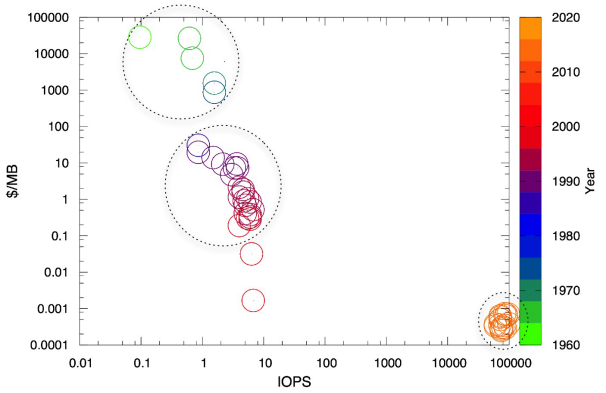

Figure 1: Storage devices over time according to cost (USD) per MB and IOPS. Data from McCallam and others.

This clustering is indicative of the evolution of the underlying niche. Storage devices exist in response to the market, which can be abstracted to a set of constraints that we claim mimics the properties of their ecological counterpart. While tracking species interaction within a given environmental niche is central to ecological modeling, Holt proposed that niches can be disentangled from their environments and, in fact, change in a Brownian manner along various axes [1]. The key attributes of the niche model are that niches are multidimensional and, at any given point, a species' placement in the niche is determined by the point where it outcompetes everything else. Within an environmental niche, different species compete for resources and eventually all of the ecological niches are satisfied or evolve. A model where we can use multi-variate regression to understand the evolution of the niche will then allow us to design a set of evolutionary algorithms modeling the competition of different devices, ultimately allowing us to better reason about the future of storage.

References

[1] R. D. Holt and M. S. Gaines. Analysis of adaptation in heterogeneous landscapes: implications for the evolution of fundamental niches. *Evolutionary Ecology*, 6(5):433–447, 1992.

[2] M. H. Kryder and C. S. Kim. After hard drives-what comes next? *IEEE Transactions on Magnetics*, 45(10):3406–3413, 2009.