# Technology Trends Favor Thick Clients for User–Carried Wireless Devices

Thad Starner
Georgia Institute of Technology

October 26, 2001

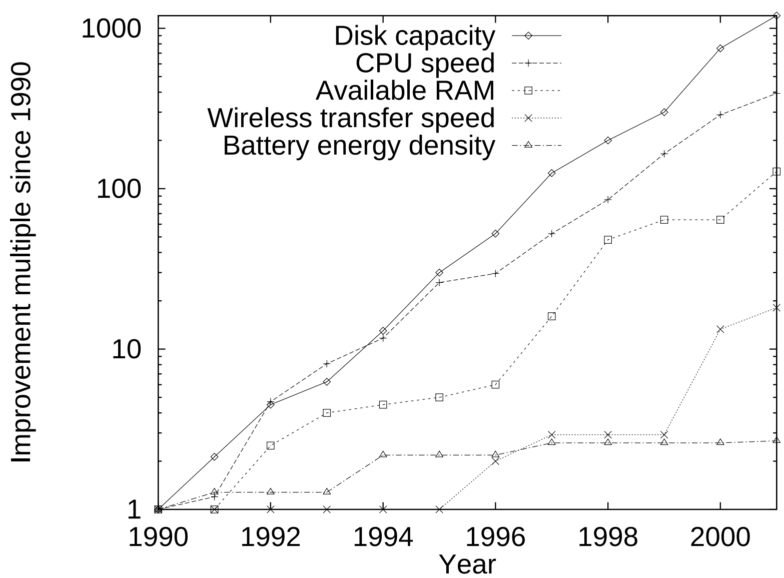


Figure 1: Laptop technology improvements since 1990.

A thin client approach to mobile computing pushes as many services as possible on a remote server. However, as will be shown, technology trends indicate that an easy route to improving thin client functionality is to "thicken" the client through addition of disk storage, CPU, and RAM. Thus, thin clients will rapidly become multi-purpose thick clients. With time, users may come to consider their mobile system as their primary general–purpose computing device, with their most used files maintained on the mobile system and with desktop systems used primarily for larger displays, keyboards, and other non–mobile interfaces.

Personal digital assistants (PDAs) and wearable computers suffer from a tension between providing more functionality and maintaining a form factor appropriate for carrying on the body. The thin client approach observes that we will not, at least in the foreseeable future, be able to pack as much technology in a pocket-sized device as can be maintained in a server in a machine room. Therefore, the thin-client approach depends on a wireless network to provide connection to a remote server for resources and provides minimal functionality on the device carried by the user. In its simplest form, the thin client would consist of a means of input and a display that is updated from the wireless network. WAP phones are a variant of this approach.

The idea of a thin client is not new — it was and, to some extent, still is used successfully for text and graphics terminals connected by wired networks to some central server. However, there are two main differences between wireline terminals and the wireless, body–worn devices being discussed today: limited power and the expectation of loss of network connectivity. The latter is easier to discuss. There are many places where providing full–speed wireless Internet coverage will remain financially infeasible for the foreseeable future. For example, as I'm writing these words I am traveling in a tunnel in the Swiss Alps. Providing wireless connectivity in a rural community with such a harsh climate is difficult enough without having to address the wireless shadows caused by Switzerland's mountains and tunnel systems. Another, more common example of wireless disconnect is during airplane flights — especially during "taxi, take-off, and landing." Providing connectivity during this time is certainly feasible but regulated against.

In its most strict form, a thin client would simply not operate during network outages. However, with the addition of local storage to cache the user's current data and a local processor to manipulate it without the network, the outage could be transparent as long as the user does not access data located outside the local cache. The more data that the thin client caches and the more processing power provided, the greater chance that the user can continue editing his presentation or playing his video game uninterrupted. However, when does a "thin–client" become a "thick–client?" When is it more profitable to think of, and to manufacture, a mobile device as its own separate entity as opposed to an interface to a network?

The progression of technology in the last decade for laptop computers, a technology now mostly mature, is shown in the figure. The graphs depict increases in performance as multiples of the state of the technology from 1990. Due to the exponential nature of the improvements, the y–axis is on a logarithmic scale. A high end machine in 1990 (the base value of 1 in the graph) was a 16 MHz 80386 with 8M of RAM and 40MB of hard drive space. Processor performance is compared in terms of Intel's iCOMP® index as derived from http://www.cpuscorecard.com; RAM and disk storage are compared by size in megabytes; wireless networks are compared on maximum nominal bits per second of data transfer; and battery energy density is determined by the type of technology used (Nickel Cadmium, Nickel metal Hydride, or Lithium Ion) and the progression these technologies have made in increasing the Joules stored per kilogram (J/kg). In general, the laptop technology represented in the graphs is that which, if repackaged in a body worn device, could be used while standing on a street corner in a major United States city. In other words, while fifteen pound "luggable" computers with more advanced technology may have been available during the years shown, only laptop computers approximately seven pounds and lighter were considered (much of the weight of these machines was due to their high-power displays and their batteries). This consideration is significant as laptop size limits the amount of heat that can be dissipated from the device — a major constraint in modern laptops.

The statistics shown were determined by examining the typical specifications of the highest–end laptop advertisements in the December issue of Byte and PC Computing magazines for each year. The wireless connectivity graph represents the author's tracking of the commercial, license–free city–wide networks available in the United States. Until the mid–90's, wireless connectivity consisted of modem access over analog cellular phones and was not reliable while moving. While the specific technology examined for the creation of the graph is often inappropriate for PDAs and wearables where battery–life and heat dissipation are limiting issues, the graphs serve to observe trends that should apply to such devices in the future.

In examining the graphs, the most striking trend is the 1200x increase of disk space demonstrated by the survey. Since 1993, disk capacity has continued to increase significantly faster than technology governed by Moore's Law as reflected in the CPU and RAM graphs. The current rate is a 2x increase in disk density every 12 months. According to IBM this rate is sustainable for at least several more generations [5]. Thus, in five years a mobile user may be carrying over a terabyte of disk storage. Not only will the user be able to store his documents on person but also his video collection!

The laggards in the graphs are wireless transfer rates and battery energy density. In particular, energy density has only increased by approximately 3x in the past decade! This result is particularly troublesome in that energy density effects both the size of the mobile device and the exploitation of the other trends.

City–wide wireless services have additional constraints besides their slow improvement. The coverage and speed of wireless networks are limited by the available hardware at the time of deployment, and the networks themselves are often controlled by companies who do not manufacture mobile devices. Thus, a device designer can not depend on the continued increase in data rates or even the availability of a wireless network, as evidenced by the recent discontinuation of the Metricom 128kbps service.

While increases in transistor and disk density and lower operating voltages continue to improve the functionality of RAM, disk, and CPU available per

Joule of energy used, wireless networks are still constrained by the inverse square power law. In practice, the amount of power needed to transfer data can increase with the fourth power of the distance to the network node [1]! Wireless service providers can combat this problem by increasing the number of wireless transponders in a given area, but average connection distance decreases slowly with the increased number of transponders (assuming an even distribution of service and service need in an area).

To provide improved thin clients, the mobile designer should concentrate on caching as can be provided by the rapid increase in disk capacity and RAM. In addition, to help alleviate power use by wireless networks, the designer can utilize increased processor capabilities and RAM to compress the data to be transferred. Once transferred, powerful local processors can allow the user to manipulate the data as much as possible with local resources. However, these tools — disk capacity, RAM, and CPU are precisely the characteristics that normally define a thick client!

In practice, the author is exploring the opposite extreme to a mobile thin client. Instead of maintaining data on a server, his data is stored on a wearable computer. The master copy of a document, including this one, is maintained and manipulated primarily on the wearable. Larger desktop machines are used mainly for computation and high–speed network access for experimentation. Significant benefits of this approach are that the user always has access to the most up–to–date data and the user maintains some sense of ownership and physical security of the data.

A potential objection to a mobile, general–purpose system with large storage capacity is the size of the device required. In actuality, given the trends discussed, the main size constraint will be the human interface. For example, keyboards can only be so small before typing speeds are affected, and displays, even head–mounted ones, become unusable if they are reduced in size beyond a limit. Given these interface constraints and the continuing issue of power, the trends above should be exploited to provide a wide variety of devices that allow the consumer to trade size for functionality, as is the case with the current laptop market.

The effort to provide computational support while mobile will involve many challenges in the coming years. Given past technology trends, mobile thick clients will play an increasingly important role in the development of the mobile marketplace, perhaps replacing the desktop as the predominant location for a user's personal files.

Pervasive computing, IEEE's newest magazine starting in 2002, will discuss such trends and challenges in mobile computing. In addition, the references below provide more information on this topic.

## References


[1] J.-H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *INFOCOM (1)*, pages 22–31, 2000.

[2] J. Kistler and M. Satyanarayanan. Disconnected operation in the coda file system. *ACM Trans. on Computer Systems*, 10(1), February 1992.

[3] M. Satyanarayanan. Pervasive computing: Vision and challenges. *IEEE Personal Communications*, August 2001.

[4] T. Starner. The challenges of wearable computing. *IEEE Micro*, July 2001.

[5] D. Thompson and J. Best. The future of magnetic data storage technology. *IBM Journal Research and Development*, 44(3):311–319, May 2000.