# Dell Dynamic Power Mode: An Introduction to Power Limits

By: Alex Shows, Client Performance Engineering

Managing system power is critical to balancing performance, battery life, and operating temperatures. In general, increasing power increases performance and provides a better user experience. But increased power also generates more waste heat and can impact battery life.

Dell Dynamic Power Mode is a set of related product features and design methodologies that control system component power. This includes power policies, software, firmware, and hardware that, in combination, **maximize performance within the constraints of the design**.

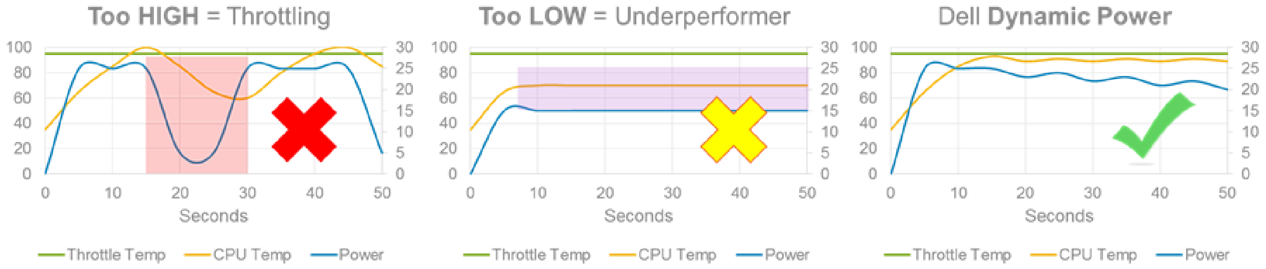

**Figure 1** – Dynamic power **maximizes performance** within temperature limits

**Background on Power**

The term "Dynamic Power" is almost self-explanatory: power is adjusted dynamically. But why and how is power adjusted? To answer these questions, we must consider the relationships between power and other factors in the system's design and operation.

Silicon-based electronics require power to function: for example to store, retrieve, and display data. Often the performance of these electronics is based on some kind of internal or external clock mechanism, and the rate of that clock corresponds to the rate of power consumed. Want higher performance? Raise the clock frequency. This is not without limits – notably, physics.

When sufficient cooling capability isn't possible due to the mechanical constraints of a design, **power limits can be used to** maintain desired temperatures.

As power is consumed and used by electronics, heat is a byproduct. Sometimes the power required by the part is so low that it can operate efficiently within its rated temperature range and not require any kind of cooling solution. Sometimes a simple heat spreader (such as copper material adhered to a chip or group of chips) is sufficient to keep the part(s) operating within their rated temperature limits. But, high power devices such as CPUs need more energy absorption and cooling capacity to function within their rated temperature limits.

When sufficient cooling capability isn't possible due to the mechanical constraints of a design, **power limits** can be used to **maintain desired temperatures**.

**Background on Cooling**

But why would we adjust power? To answer that requires a brief explanation of cooling solution design and operation.

When a cooling mechanism like a heatsink is attached to a chip, this slows the increase in temperature as the heatsink absorbs the heat generated. But, it only slows the temperature increase of the chip, and at some point can reach a point called thermal soak where the cooler is no longer effective at absorbing energy from the chip.

Chips will typically be rated for certain power, and in the case of the CPU, this is the Thermal Design Point (TDP). GPUs have similar thermal design points, and while other parts in the system have these design points, the focus for Dell Dynamic Power is on the CPU and GPU, as these are two of the highest power components in the system.

CPU and GPU manufacturers will often provide guidance on the thermal characteristics of their chips, including nominal ratings. So long as an implementation can supply sufficient nominal power, and provide adequate cooling so that the chip never exceeds a particular threshold (typically maximum thermal junction temperature, or $T_j$Max), the chip manufacturer guarantees a certain operating frequency can be maintained by the chip. Maintaining a nominal operating frequency is important for adequate performance, but what if we want more performance?

**CPU and GPU Performance**

Heatsinks help to slow the rise in temperature as input power increases. The larger the heatsink, the more energy that can be stored, and heatsinks with fans and lots of surface area can transfer heat into the environment and away from hot components. Thus, temperature rises and falls with power, albeit delayed as the heatsink soaks up heat and distributes it elsewhere.

CPU and GPU manufacturers exploit this behavior to achieve greater performance. So long as sufficient power and cooling are provided, the CPU or GPU can run at higher frequencies than their nominal rating for short periods of time. These short periods can't be sustained though, because the heat generated will eventually exceed the rating of the part if the thermal solution (e.g. heatsink and fan) is designed for some lower, nominal power rating.

For short periods of time a part may be able to run at a higher frequency than nominal, and this is marketed as "Turbo" (Intel® Core™ processors) or "Boost" (NVIDIA® GeForce™ graphics processors). This is a very useful behavior to exploit in client platforms like notebooks, because the most common usage model on these systems is interactive work. That is, one might open an application, wait for it to load, draft an email, read another email, open a web page, and finally read an article on the page. There are periods of time during this kind of use case that the system is idle followed by periods of activity. It is exactly this "burst" usage that creates the opportunity to provide significantly higher power than nominal, because of the idle periods in between bursts when the CPU and/or GPU have time to cool.

To summarize, higher performance can be delivered for short periods of time by delivering higher power, so long as temperature limits are not exceeded. CPUs and GPUs have internal mechanisms to regulate power and temperature by adjusting operating frequency to stay within defined limits.

**Why Dynamic Power?**

If the CPU and GPU have internal mechanisms to control temperature and power based on frequency, why use Dynamic Power? Why wouldn't we simply raise all the power limits very high and let the parts regulate themselves?

While certain PC designs like desktop systems might include heatsinks which are so large that they never reach their thermal limits *ever*, notebook systems face a much greater challenge: thinness.

With each generation, notebook PCs are increasingly thinner. The thinner the system, the closer all those hot components get to the surface and the less space there is to provide cooling. One of the effects of thinning the design is that a component such as a CPU or GPU may no longer be allowed to run to its peak operating temperature. If the upper-temperature limit of a CPU or GPU is, let's say, 100 degrees Celsius (212 degrees Fahrenheit), any surface very close to the heatsink connected to that device will be too hot to handle! In fact, Underwriter's Laboratory (UL) sets device safety standards for surface temperatures which are *much* lower than the maximum safe operating temperatures of CPUs and GPUs.

Therefore, in many PC designs, high power components such as CPU and GPU cannot be allowed to run to their thermal limits, and must be constrained. We can constrain temperature directly by reducing their operating frequency, or indirectly by placing power limits on the parts and allowing them to adjust frequency to stay within those power limits.

> Setting a fixed power limit based on steady-state thermal soak **prevents the system from achieving its full potential.**

Some designs might choose to simply characterize the system under heavy load, find the temperature of each component which corresponds to a skin temperature limit, and then restrict that part to that lower power level. But, setting a fixed power limit based on steady-state thermal soak prevents the system from achieving its full potential, because for all the time spent before thermal soak, higher power could deliver higher performance *without* exceeding temperature limits.

Figure 2 is a simplified illustration of power and temperature with fixed limits applied, based on worst-case thermals. After a long time running at high power, the component will eventually reach thermal soak so long as that temperature is within its operating range. In this example the temperature at thermal soak is below the throttle temperature limit for the platform, but for all the time prior to thermal soak, the CPU underperforms due to the power constraint.

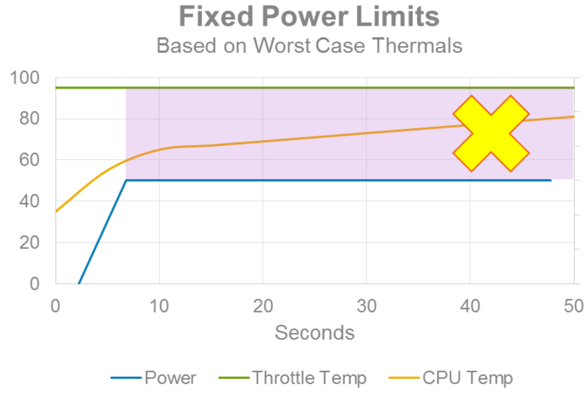

**Figure 2** – Fixed Power

Spending any amount of time limited by some future possible temperature value is bad for performance. The lower the worst-case thermal soak power limit, the lower the performance when thermals are not a constraint.

### Dell Dynamic Power Mode

Every platform design is different, and to achieve maximum performance in designs in which components are constrained by skin temperatures, real-time measurement of device temperatures is critical.

Rather than "set and forget" using worst-case thermal conditions, power limits can be adjusted dynamically based on measured thermal data. In order to do this accurately, quite a number of thermal measurements must be taken in Dell's laboratories. Each system design is instrumented with thermocouples which monitor skin temperatures at a dozen or more locations on the system.

Data logs of skin temperature readings are analyzed and correlations are established between component power draw, operating frequency, device temperature, skin temperature, and other important factors. This set of correlations enable Dell platforms implementing Dell Dynamic Power Mode to predict the temperature of the *surface* of the system using only *internal* device measurements.

The results of this analysis is a matrix of power limits and "thermal rules" which can be applied in real-time to adjust power limits. These rules for adjusting power in response to temperature can then be implemented via an Embedded Controller (EC) connected to each component, or via software running as a service on the operating system.

Dell Dynamic Power Mode implements dynamic power limits through a combination of EC and OS-based controls, ensuring that each platform which incorporates this feature delivers maximum performance within its skin temperature constraints.

One such software service on PCs with Intel® Core™ processors is the Intel® Dynamic Platform Thermal Framework (DPTF), which provides a framework that is a natural fit for dynamic power policies defined and implemented by Dell.

Contrast Figure 3 below with the fixed power limits shown previously. Dell Dynamic Power maximizes performance by raising the power limit as high as possible without exceeding the temperature limit.

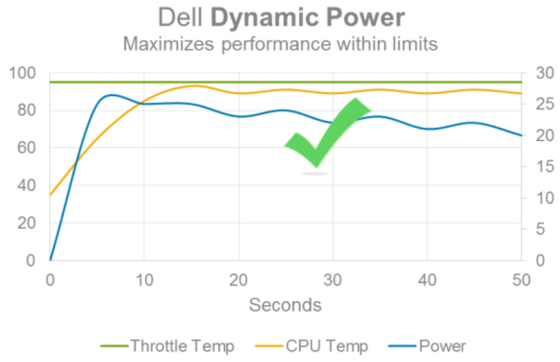

**Figure 3** – Dell Dynamic Power behavior

Notice how the blue line (power) is smoothly adjusted up and down in a sinusoidal pattern in order to keep the temperature from reaching the thermal limit. In this hypothetical example, that limit is a throttle temperature. But, dynamic power limits mean that even the temperature thresholds for power reduction themselves can be dynamically adjusted.

One can easily see how this could be applied to set user-preferred temperatures for their environment. While docked, the system might set a higher limit than if the system were held or worn. In either case, the *temperature limit* can be adjusted in real-time via the Dell Dynamic Power Mode policy to deliver the maximum performance possible within that adjusted constraint.

## Conclusion

Achieving higher system performance often requires increased power consumption. But, increasing power consumption creates more heat. Further, many designs such as thin notebook PCs may be constrained by skin temperatures. Simple approaches to thermal management might limit power to stay within worst-case thermal limits but fail to maximize performance.

Dynamic power limits use a platform-specific set of rules, based on controlled laboratory measurements of internal and external temperatures, to predict system surface temperatures based on internal readings. By adjusting power in real-time using predicted surface temperatures, Dell Dynamic Power Mode **maximizes performance within the constraints of the design**.