August 20, 2018

Detecting Heatsink Types for Socketed Processors

Scott Faasse
Hewlett Packard Enterprise

Kevin M. Cash
Hewlett Packard Enterprise Show/hide Remove author Email First

Detecting Heatsink Types for Socketed Processors

Abstract

Servers often are designed to support a socketed processor in order to allow the end user the ability to customize their compute solution for their needs. Server processor vendors offer a variety of different processor models that can be installed. These various processor models can have differing technical specifications that include core count, cache size, operating frequency limitations, memory capacity, as well as power and thermal cooling requirements. An individual server design could easily support a range of processor models from those that have few cores, providing low performance, all the way up to dozens of cores that providing high performance. Since compute resources such as cores consume power, the range of processor power can be as wide as 200W. To keep costs of a total solution down, different heatsinks are often designed and sold to support this wide range: cheap extruded aluminium heatsinks for lower performing CPUs and expensive heat piped heatsinks for higher performing CPUs.

Overview

A problem with heatsink and platform designs that persists even today, is that system firmware, which is responsible for thermal, power, and performance management cannot readily determine which type of heatsink is installed. System firmware can easily identify processor model, but not the heatsink that is installed on top of it. It can only be assumed that manufacturing or service personnel installed the proper heatsink for a given processor model. The lack of ability to identify the heatsink installed could lead to issues if a mismatch occurred for which system firmware could not actively address or notify. In addition, if a customer were to install a higher performing (more expensive) heatsink onto a less constrained processor model, system firmware is unable to take advantage of this configuration to slow fans down in order to be more efficient or allow the processor to run at a higher performing level.

Description

There are two heatsink and board design modifications that we've identified that would allow system firmware the ability to detect the type of heatsink that may be installed with the processor.

Method 1: Pogo pins are mounted on the heatsink (see figures in attached illustrations) in a pattern that is unique to each heatsink type. Depending on where these pins contact corresponding pads on the motherboard, on-board circuity (pullups, buffers, GPIOs, etc.)

would allow system firmware to identify which heatsink is installed. The number of pin/pads (n) used provides 2^n possible detectable heat sink types (3 pin/pads = 8 possibly types).

Method 2: Somewhat the inverse of first method, rather than installing protruding pins on the heatsink surface, the edge of the heatsink plate is recessed via CNC (much like the bolt hole pattern on existing heatsink designs.) Spring-pin style connectors are installed on the motherboard which will force contact for non-recessed areas of the heatsink plate. Similar circuitry on the motherboard described in method 1 are used so that system firmware can detect the type of heatsink installed.

Upon detection of the heatsink type, the system firmware can optimize thermal, power, and performance management policies in order to take advantage of better power efficiency and/or compute performance headroom.

An additional benefit in being able to detect heatsink type is the ability to detect if unsupported (3rd party, competitor, older generation, wrong model, etc.) heatsinks are installed. The placement of the pins/recesses for both methods could be designed to be unique to an original equipment manufacturer (OEM). An unsupported heatsink could be detected and the system prevented from behaving optimally or even booting.

Figure 1 – A heatsink using method 1 which has two points of contact to motherboard's mating leaf spring connector ("two feet")

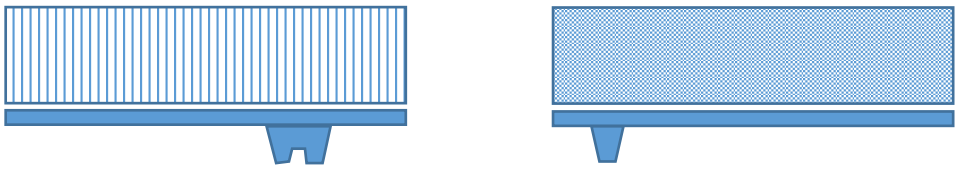

Figure 2 – different supporting heatsink with just one contact point ("one foot")

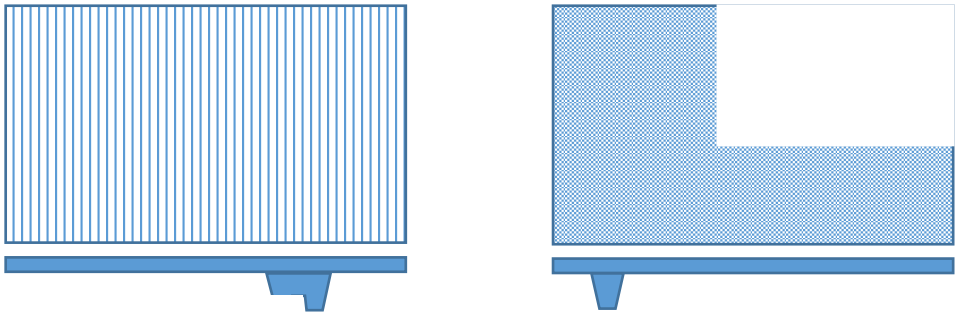

When the heatsink is installed, its “padded feet” depress on motherboard leaf spring that closes the circuit. Whether the circuit is open or closed can be read by system firmware in order to determine if this particular heatsink has been installed. Shown below is a close up of the above fig.2 “one foot” coming in contact with motherboard’s switch.

Figure 3 – shows how heatsink’s “one foot” depresses motherboard switch to close the circuit.

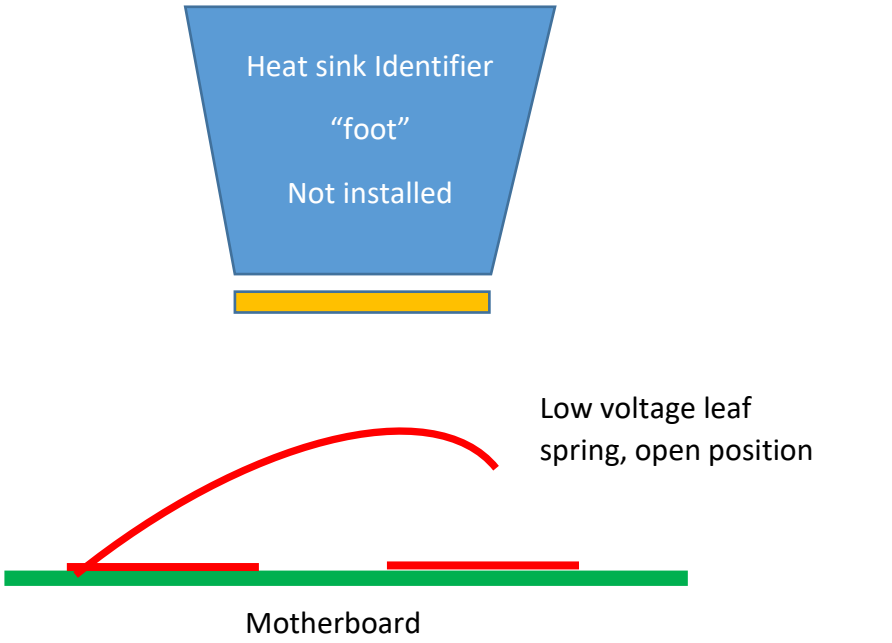

Disclosed by Scott Faasse and Kevin M. Cash – Hewlett Packard Enterprise