Technical Disclosure Commons

Defensive Publications Series

March 21, 2019

Throttling downlink throughput to mitigate device temperature increase

Shivank Nayak

Blake Kragten

Siddharth Ray

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Throttling downlink throughput to mitigate device temperature increase

ABSTRACT

The temperature of a mobile device can increase due to heavy use, e.g., high-speed downloads, large computational load, etc. Sustained periods of high temperature can damage the mobile device. The techniques of this disclosure reduce downlink throughput upon detection of device temperature that exceeds a threshold. Throughput is reduced, e.g., by signaling the thermal state to the network, by reporting lower channel quality indicator (CQI) values to the network, etc. After the temperature drops to a safe level, throughput is brought back up in a phased manner.

KEYWORDS

device temperature; thermal state; channel quality indicator (CQI); downlink throughput; attach procedure; connected mode signaling; 3G; LTE; 4G; 5G NR; tracking area update; data throttle; throttling

BACKGROUND

The temperature of a mobile device can increase due to heavy use, e.g., high-speed downloads, large CPU/GPU load, etc.; due to environmental factors, e.g., being positioned on a car dashboard in the sunlight; etc. Sustained periods of high temperature can damage the mobile device hardware.

DESCRIPTION

The techniques of this disclosure reduce downlink throughput upon the detection of device temperature that exceeds a threshold. Throughput is reduced e.g., by signaling the thermal state to the network, by reporting lower channel quality indicator (CQI) values to the network, etc. After device temperature reduces to a safe level, throughput is brought back up in a phased

manner. Per the techniques, temperature can be mitigated using a user equipment (UE) centric

approach, a network-centric approach, or a hybrid approach.

UE-centric temperature mitigation

In a UE-centric approach, temperature mitigation is primarily driven by the device that is
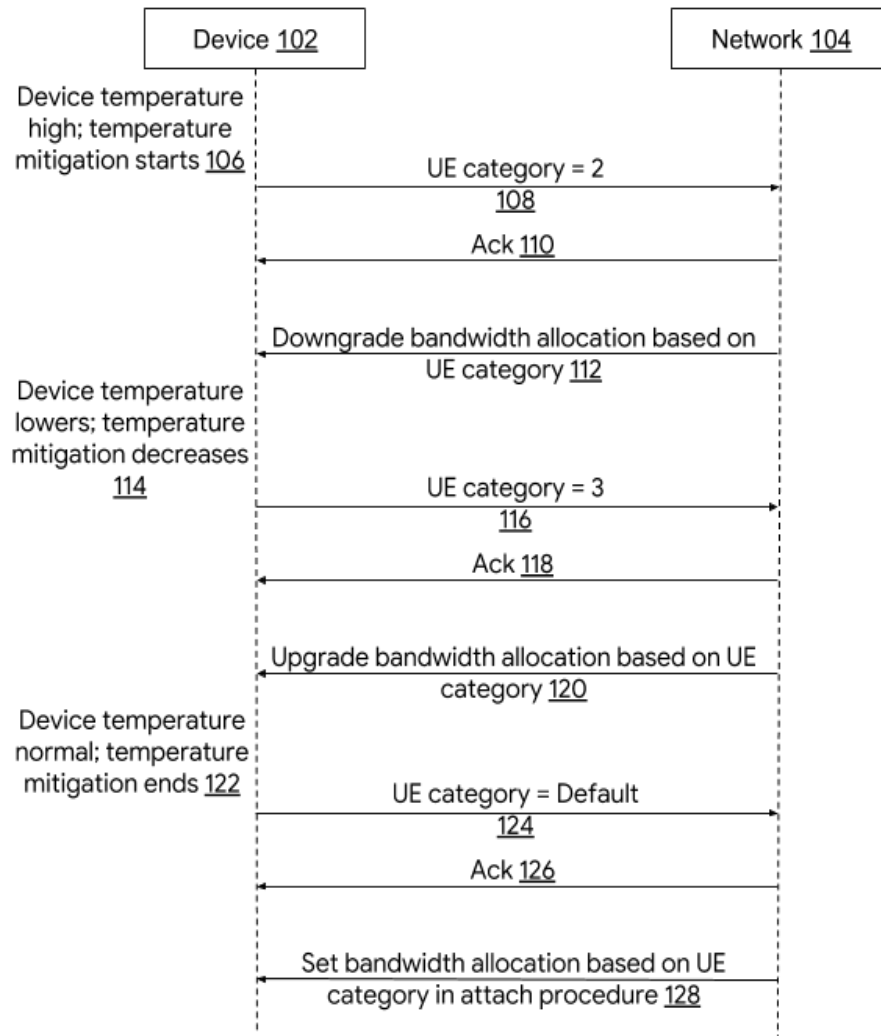
undergoing temperature change.

Fig, 1: Temperature mitigation using UE category to set downlink bandwidth allocation

Fig. 1 illustrates an example of signaling between a device (102) and a network (104) to

mitigate temperature change. Wireless standards, e.g., 3GPP [1], define bandwidth categories for

a UE. The category of a UE establishes the data throughput of the UE. The UE indicates its highest UE category during an attach procedure, and changes its category during connected mode based on measurement of its own temperature. After the device detects a high temperature (106), e.g., above a static or dynamic threshold or approaching that threshold at a positive rate of change, temperature mitigation starts. The device signals a relatively low bandwidth category (108), e.g., a UE category 2, to the network using connected mode signaling. The network responds with an acknowledgement (110) and downgrades bandwidth allocation (112) based on the received UE category.

After the device temperature drops (114), e.g., at a certain negative rate of change or to a different static or dynamic threshold, temperature mitigation decreases. The device signals a higher bandwidth category (116), e.g., a UE category 3, to the network using connected mode signaling. The network responds with an acknowledgement (118) and upgrades bandwidth allocation (120) based on the received UE category.

After the device temperature normalizes (122), e.g., to a zero rate of change and/or below a static or dynamic threshold, temperature mitigation ceases. The device signals a default UE category (124) to the network using connected mode signaling. The network responds with an acknowledgement (126) and sets bandwidth allocation (128) based on the UE category in the attach procedure
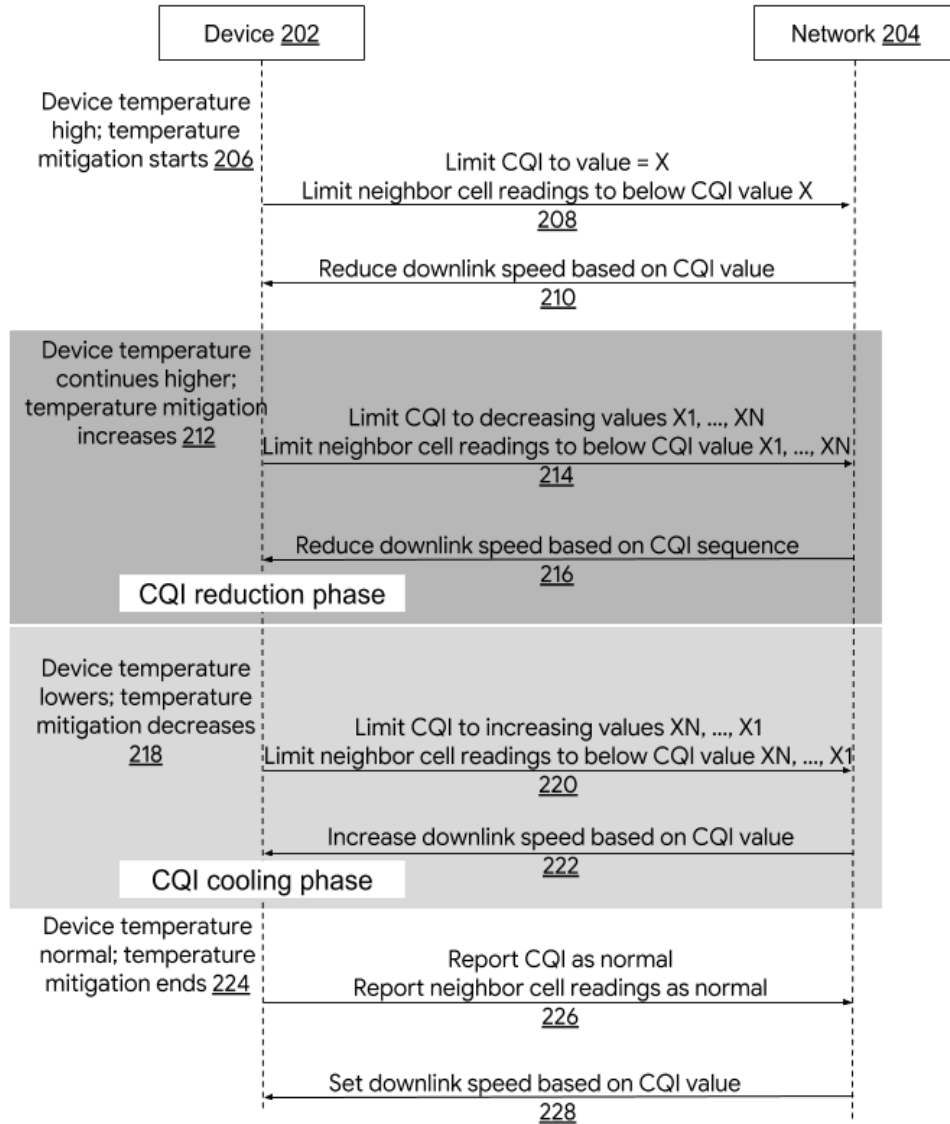
Fig. 2: Temperature mitigation using CQI to set downlink throughput

Fig. 2 illustrates an example of temperature mitigation using channel quality indicator (CQI), which achieves a more granular control of downlink throughput. CQI is an indicator of the capacity of the communication channel. CQI is typically fed back by the device to the network so that the network can set the downlink speed based on channel quality.

Per the techniques of this disclosure, the device mitigates temperature increase by gradually decreasing reported CQI which in turn causes the network to decrease downlink

throughput. A device (202) in communication with a network (204) detects a rise in temperature and starts temperature mitigation (206). The device signals a reduced CQI (X) to the network (208). To prevent flip-flopping across cells or the network handing over the device to a neighbor cell, the device also reports neighbor-cell CQI readings as being below the CQI value X (208), or altogether disables neighbor-cell search. The network reduces downlink speed to the device based on the received CQI value (210).

If device temperature continues higher (212), this prompts further temperature mitigation; this is referred to as the CQI reduction phase. The device signals a sequence of generally decreasing CQI values X1, …, XN to the network (214). With each CQI value Xi, the device reports neighbor-cell CQI readings as being below the CQI value Xi (214), or altogether disables neighbor-cell search. The network reduces downlink speed to the device based on the sequence of received CQI values (216). In this CQI reduction phase, the CQI is gradually stepped down, e.g., until the minimum CQI is reached. Stepping down the CQI value occurs when specific temperature thresholds are reached. Within the device, each CQI value has a configurable temperature set point, which lowers the CQI to the corresponding value, and a temperature clear point, which increases the CQI to the next higher value.

At a sufficiently low downlink speed, the temperature of the device starts reducing (218), and the device enters a CQI cooling phase to decrease temperature mitigation. The device signals a sequence of generally increasing CQI values XN, …, X1 to the network (220). With each CQI value Xi, the device reports neighbor-cell CQI readings as being below the CQI value Xi (220), or altogether disables neighbor-cell search. The network increases downlink speed to the device based on the sequence of received CQI values (222).

In this CQI cooling phase, if the device starts cooling and hits the temperature clear point for a CQI value, the CQI value goes back to the immediately higher CQI value. The CQI cooling phase ends when the initial mitigation-clear temperature is reached. If at any point during the CQI cooling phase the temperature starts increasing again and hits a CQI step-down temperature set point, the device reduces the reported CQI value, thereby starting the CQI reduction phase again.

When the device temperature normalizes (224), the device reports CQI and neighbor-cell CQI as normal (226) in order to complete temperature mitigation. Throughout this process, the network sets downlink speed based on the received CQI value (228).
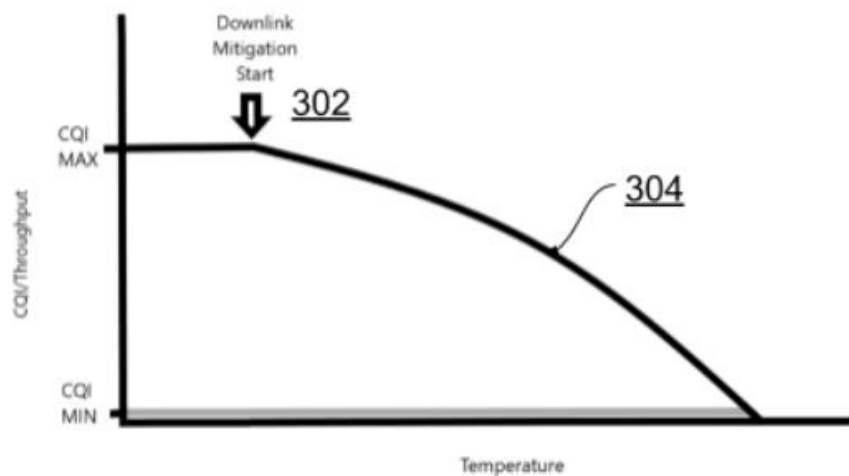

Fig. 3: Throughput against temperature

Fig. 3 illustrates the reported CQI, and hence throughput, as a function of temperature. As temperature rises with time, it hits a threshold that prompts temperature mitigation (302). After this point, the increasing temperature causes the device to lower the reported CQI (throughput) (304). The process is reversed as the temperature reaches maximum and then reduces.

In this manner, smooth reduction in data throughput effects a smooth reduction in temperature.

Network-centric temperature mitigation

In network-centric temperature mitigation, the device communicates to the network its thermal abilities, e.g., support for thermal throttling, during the attach/ tracking-area update procedures. During connected mode, the device periodically communicates its temperature to the network. Upon receiving temperature information, the network throttles downlink data to achieve temperature reduction at the device.
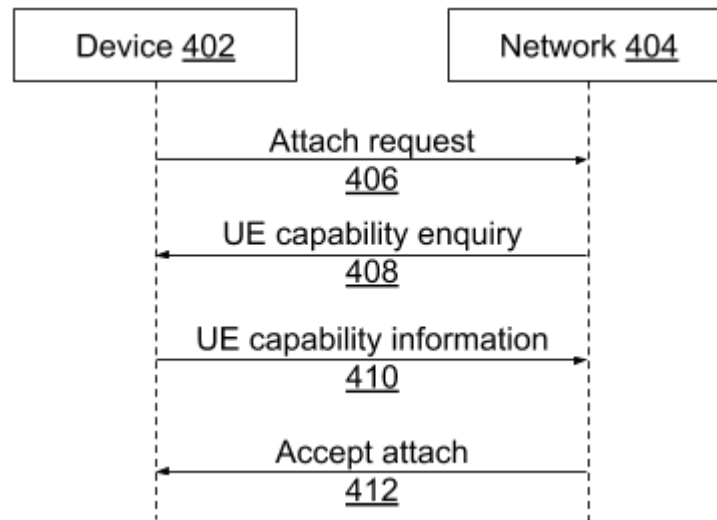

Fig. 4: Device-network messaging to communicate UE capability

Fig. 4 illustrates an example of device-network messaging to communicate UE capability during attach/ tracking-area update procedures, per techniques of this disclosure. The device (402) sends an attach request message (406) to the network (404). In return, the network sends an enquiry relating to the capabilities of the device (408). The device responds with a UE capability information message (410). In addition to the traditional information carried by the UE capability message, e.g., supported wireless bands, carrier aggregation combinations, etc., the UE capability message also carries information relating to the thermal abilities of the device. In particular, thermal abilities include the maximum supported temperature for various device

regions. The network accepts the attach request (412), and communicates to the device its support for thermal throttling in the attach / tracking-area update accept message. This device-network message exchange enables the device to periodically communicate its thermal state to the network when in connected mode.


Fig. 5: Device-network messaging to mitigate temperature in a network-centric approach

Fig. 5 illustrates device-network messaging to mitigate temperature in a network-centric approach. The device (502) is in communication with the network (504) in connected mode. Upon detecting a high temperature (506), the device sends a message (508) to the network, e.g., "Skin temperature: 55 ^{0}C," using connected mode signaling. The network acknowledges receipt of the message (510) and reduces downlink speed (512) to a level commensurate with device temperature.

Fig. 6: Device-network messaging to mitigate temperature in a network-centric approach

Fig. 6 illustrates another example of device-network messaging to mitigate temperature in a network-centric approach. In this example, the network offloads the device to a lower rate technology in order to reduce device temperature. The device (602) is in communication with the network (604) in connected mode. Upon detecting a high temperature (606), the device sends a message (608) to the network, e.g., "Skin temperature: 55 ^{0}C; CPU temperature: 70 ^{0}C," using connected mode signaling. The network acknowledges receipt of the message (610) and offloads the device to a lower data-rate technology (612), e.g., from LTE to WCDMA, from 5G NR to LTE, from cellular to WiFi, etc.

Hybrid approach to temperature mitigation

In a hybrid approach, network-centric and UE-centric mitigation co-exist such that UE-centric mitigation is used as fallback if the temperature reduction provided by network-centric mitigation is insufficient. In the hybrid approach, in addition to the thermal level set to start notifying the network of thermal conditions, the device sets a self-mitigation thermal level for each type of thermal state. If the self-mitigation level is reached, the UE indicates to the network a reduced data rate or UE category to force throughput reduction.

Various thermal states can be used as measures of temperature, e.g., one or more of: skin (device surface) temperature, CPU temperature, junction temperature, battery temperature, etc.

Alternatively, in order to reduce temperature, an already connected UE can perform another attach with reduced UE capability. Still alternatively, a temperature-specific signaling message between device and network can be incorporated into the standard to limit downlink throughput.

CONCLUSION

The techniques of this disclosure reduce downlink throughput upon detection of device temperature that exceeds a threshold. Throughput is reduced, e.g., by signaling the thermal state to the network, by reporting lower channel quality indicator (CQI) values to the network, etc. After the temperature drops to a safe level, throughput is brought back up in a phased manner. In this manner, temperature-related damage to a mobile device or loss in service is prevented.

REFERENCES

[1] 3GPP, Section 36.606 Table 4.1-1. 3rd Generation Partnership Project.