

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 May 2009 (14.05.2009)

PCT

(10) International Publication Number
WO 2009/062186 A2

(51) International Patent Classification:
H04W 52/02 (2009.01) *H04B 1/16* (2006.01)

Robert [AU/AU]; 15 Cassandra Avenue, St Ives, NSW, 2075 (AU).

(21) International Application Number:
PCT/US2008/083050

(74) Agent: **VINCENT, Lester, J.**; Blakely, Sokoloff, Taylor & Zafman LLP, 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (US).

(22) International Filing Date:
10 November 2008 (10.11.2008)

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/983,563 9 November 2007 (09.11.2007) US

(71) Applicant (*for all designated States except US*): **G2 MICROSYSTEMS PTY. LTD** [AU/AU]; Level 15, 24 Campbell Street, Haymarket, New South Wales, 2000 (AU).

(71) Applicant (*for NA only*): **LINDEEN, Gordon, R. III** [US/US]; 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (US).

(72) Inventors; and

(75) Inventors/Applicants (*for US only*): **HOGAN, Michael, Thomas** [AU/AU]; 6 Kedron Avenue, Beecroft, NSW 2119 (AU). **RYAN, Phillip, John** [AU/AU]; 42 Percival Road, Stanmore, NSW (AU). **CRAWFORD, Jason,**

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report*

(54) Title: EARLY RECEIVER SHUTDOWN

(57) Abstract: A receiver of a mobile device is deactivated and reactivated between frames of a data transmission system. In one example, a radio receiver of a mobile device is activated. The radio receiver listens for preambles of frames transmitted by radio transmitters and receives a frame preamble. The receiver is deactivated upon receiving a frame preamble. The receiver is then reactivated after a time determined based on the received frame preamble to receive a subsequent frame preamble. The receiver receives the subsequent frame preamble, determines an address for the subsequent frame, and receives the subsequent frame, if the address corresponds to the mobile device.



WO 2009/062186 A2

EARLY RECEIVER SHUTDOWN

BACKGROUND

Field

The present description pertains to the field of power consumption for mobile radio devices in radio frame networks and, in particular, to deactivating and activating a receiver of such a device between receiving frames in the network.

Related Art

Wireless network devices are increasingly popular for a wide range of different applications, from personal communications, to network appliances, to surveillance and distributed control systems, to location determination and tracking systems, to home and office automation systems, to toys. Many of these devices are portable and powered by batteries or other low or limited range power sources. In a low power portable device, the wireless network features can consume a significant part of the device's total available power.

Wireless network communication protocols, such as IEEE 802.11 (a standard for wireless communications promulgated by the Institute of Electrical and Electronics Engineers) provide a protocol that a portable device can use to discover any access points (AP) or stations (STA) that are within range and to communicate data across the wireless network. Such a protocol can also allow for a portal into wired networks of various types.

In establishing network connectivity, the portable device selects an appropriate access point for its network communications. There are a variety of different approaches including selecting one that is configured by a user. Another approach is to select the access point with the highest RSSI (Received Signal Strength Indication).

If there are a large number of nearby access points, that is the access point density is high, selecting the best access point, using any approach, can be time consuming. In a typical 802.11 protocol, the mobile device sends out a probe using its transmitter. The probe is received by all of the access points within range. Each access point can then

decide whether to respond. Normally all of the access points within range will respond. The mobile device waits while the access points respond. According to the 802.11 protocol, the mobile device must then acknowledge each response. Normally, there are three or more different radio channels that an access point can use, so the mobile device repeats this process on each channel. This allows the best access point and the best channel to be discovered.

The interval between each probe response received by the mobile device is typically 1ms (millisecond), so if the mobile device can “see” a large number of access points (e.g. 30 per channel), then the mobile device must spend approximately 30ms (1 ms x 30 AP's) transmitting and receiving messages for each channel. Usually the mobile device will repeat this process for each of the three channels. In addition, in order to obtain a reliable indication of an access point's RSSI, multiple RSSI measurements are normally needed for each channel (e.g. 3), so the mobile device might be transmitting or receiving for almost 2 seconds for each location (1 ms x 30 APs x 3 channels x 3 measurements). This is before the AP can choose an AP with which to connect or, for location systems, determine the position of the mobile device or transmit any information about its position to an outside tracking system.

Each millisecond that the mobile device must stay active consumes energy for a battery-powered device, the process of selecting an access point can consume significant battery power. This is particularly true if the mobile device must repeat this process each time that it has to report its position. Similarly, remaining powered in order to receive frames from the access point will also consume power over time. This power consumption will eventually discharge a mobile device's battery. However, a longer battery life will significantly reduce the cost of using and maintaining a large number of RFID tags, or any of a variety of other wireless devices.

SUMMARY

A receiver of a mobile device is deactivated and reactivated between frames of a data transmission system. In one example, a radio receiver of a mobile device is activated. The radio receiver listens for preambles of frames transmitted by radio transmitters and receives a frame preamble. The receiver is deactivated upon receiving a

frame preamble. The receiver is then reactivated after a time determined based on the received frame preamble to receive a subsequent frame preamble. The receiver receives the subsequent frame preamble, determines an address for the subsequent frame, and receives the subsequent frame, if the address corresponds to the mobile device.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention may be more fully appreciated in connection with the following detailed description taken in conjunction with the accompanying drawings, in which like reference numerals refer to corresponding parts throughout the several views of the drawings, and in which:

Figure 1 is a block diagram of an wireless network device in an environment including multiple access points for communication, according to an embodiment of the invention;

Figure 2A is a diagram of a wireless communication frame suitable for use with embodiments of the present invention;

Figure 2B is a diagram of a PSDU portion of the wireless communication frame of Figure 2A;

Figure 3A is a block diagram of a mobile device, according to an embodiment of the invention;

Figure 3B is a more detailed block diagram of a receiver of the mobile device of Figure 3A; and

Figure 4 is a process flow diagram of establishing communications with one of the access points, according to an embodiment of the invention.

DETAILED DESCRIPTION

In a typical wireless network, there is often a process of discovery that mobile devices use. This process typically includes measuring the RSSI (Received Signal Strength Indicator) for several different APs (Access Points) and then either selecting one as a network access node, or using the measurement in a location system. For location, the best location data is obtained from APs (Access Points) that are close to the mobile

radio. For network communications, the closest APs also typically have the clearest signal, provide the highest signal strength, and require the least amount of transmit power

In one example of the invention, when scanning for RSSI values, there is no need to keep the receiver active longer than necessary to measure the RSSI. A typical frame is much longer than necessary to make the measurement and indicates the sender of the frame in the preamble. In addition, the frame may also contain enough information from which to determine how long to wait for the start of the next frame. In 802.11, for example, the LENGTH and RATE fields in conjunction with the DURATION field of the received frame can be used for this purpose.

As a result, after the RSSI measurement is made, the receiver can be deactivated until the next frame. The deactivation times can be used to save power. In one example, the receiver will listen to a pre-selected or other appropriate wireless channel until it finds a frame. As soon as a good frame is found, the preamble of the frame is received and analyzed. For an 802.11 example, the preamble will contain a sequence field, including the sequence ID, a frame type, and a from-address, indicating the sender of the frame. This data can all be written to memory for use in managing the process, among other things. At the same time, the RSSI can be measured and the measured value also written to memory. At this point, having received the preamble the receiver can be shut down for the rest of the frame.

When, according to a timer or a calculation, it is time for the next frame to be transmitted, the receiver is activated or awakened and the process can be repeated. For long frames, this allows the receiver to be shut down for many milliseconds. In some cases, it may be preferred to receive an entire frame, in order to receive error correction or detection sequences. In 802.11 sequences, the FCS (Frame Check Sum) is at the end of the frame and allows the preamble data to be checked for accuracy.

Further power savings can be obtained by staying active or awake only for frames of interest. For example, the FROM-DS field can be used to determine whether a frame is from an AP. If not, then the receiver can be deactivated until the next frame. Similarly, the from-address can be used to ignore frames for far away or otherwise unimportant APs.

Embodiments of the present invention allow power consumption to be reduced in circumstances in which the RSSI or some other quality parameter of received radio frames is to be measured. The power saving is achieved by turning off a radio, and possible also a baseband PHY (Physical layer) processor whenever possible, that is whenever these components will be inactive to turn off long enough to provide a significant power savings.

Power is saved by listening to only as much of a packet as is necessary to perform the desired function. The functions, according to some embodiments of the invention include 1) measure the RSSI, 2) classify the frame (who is it from, who is it to, is it from an AP, etc), and 3) extract protocol timing information from the frame. The protocol timing allows the receiver to calculate how much longer the current frame will be on the air. With this information, the radio can be turned off until the next frame.

These operations are useful in several different scenarios. For example, when a mobile device moves into a new area, there may be a discovery process. The mobile device tries to discover an Access Point with which to associate. Measuring RSSI or some other quality measures provides data that the mobile device can use to select an AP. Alternatively, in a location determination function RSSI can be used for location determination. The mobile device can be configured to report on the measured RSSI of several different access points. The report can then be used by a location server or other device to determine the location of the device. For autonomous location determination, the measured RSSI can be used by the device to determine its location relative to the APs.

Figure 1 is a block diagram of an example environment for a mobile device in a wireless network. The mobile device is described in more detail in the context of Figure 3. In the described example, the mobile device is a battery powered active mobile device capable of sending, receiving, and processing. The particular configuration of the mobile device is not important to the invention. Instead of a mobile device, a more complex device can be used that performs other functions beyond those described here. The mobile device can be one that is capable of receiving commands, reporting information back to the network, streaming video or audio, presenting data on an interface, determining its position, or any of a variety of other tasks. Alternatively, the mobile

device may be very simple and very low power such as a radio frequency identification smart tag.

Figure 1 also shows four APs 103, 105, 107, 109. Each AP can receive a probe from the mobile device and send a response back to the mobile device. The particular format of the probe and response is not important to the invention. A wide range of different techniques may be used. The examples described herein are based on 802.11 protocols, but any other signaling system may be used. The APs may be conventional wireless access points used in personal computer networking, or they may be specifically adapted to this or a variety of other applications.

One of the APs, AP4 109 also receives a location message. This indicates that the mobile device has selected this AP as the one with which to communicate and provides information about the mobile devices location. The selection can be done using the approach described below or any number of other approaches. The location message may contain nothing more than an identification of the mobile device, such as a mobile ID, or a MAC (Media Access Control) address. This permits the AP to establish network connectivity for the mobile device. As an alternative to a location message and connection request may be sent instead.

In some embodiments with location services, the AP can also attach a location and a time to the mobile device identity so that it can be tracked. Additional information may also be added depending on the particular application. The message may also include other information, including a location determined by the mobile device, environmental measurements by the mobile device, logged data that the mobile device has accumulated since its last message or since any other point in time, and more. The content and the configuration of the message can be selected to suit the particular articles being tracked and the environment in which they are being tracked.

AP4 is also shown as being connected through a network 111 to a personal computer 113 and a location server 115. The other APs can also be connected to this or another network, either directly, like AP4 or through other APs. These connections are not shown in order to simplify the drawing. The personal computer may be used for WAN (Wide Area Network) access, system management, communications, location tracking and also to make configuration and other changes. The personal computer or

user terminal provides access to the APs and the network for configuration, maintenance, management or any other purpose.

The location server determines the location of each mobile device. It usually stores a location map of each access point, or alternately it can store a map of RSSI values and the places where these RSSI values may be found. The location server also maintains a record of the location of each of the mobile devices that report into an AP. This record may be sent to an inventory maintenance and tracking system, a delivery fulfillment system, a customer location system or any of a variety of other systems. As an alternative, the location system can determine the location of the mobile device. It may do this based on determinations from the APs or it may receive only measurements from the APs.

In one example, the location server receives RSSI measurements from each AP that detects the probe or another signal from the mobile device. The location server 115, is coupled to an RSSI database 117. This database has a table of RSSI values for each AP. Using the table and the measurements from the AP, the location server can determine the location of the AP. This value can be stored in the location server, and, as mentioned above, sent to other devices that desire the information.

There are a variety of different ways to determine a mobile devices location including receiving GPS coordinates from the device. The location services are provided as an example of an application that may be provided with the mobile device 101. In some embodiments of the present invention, neither the location services application nor the location server, nor the RSSI table are necessary.

Figure 2A shows an example of an 802.11 wireless networking frame suitable for use in the described embodiments. While this description is presented in the context of 802.11, a variety of other wireless networking protocols may be used as alternatives. A frame such as that shown in Figure 2 may be used, for example, in the wireless link between a mobile device, such as a tag, or any other mobile device and a wireless access point (WAP) or other communication node or base station. In accordance with conventions for 802.11, a frame 207 begins with a PLCP (Physical Layer Convergence Protocol) preamble 209, followed by a SIGNAL 217, followed by the DATA section 219.

The three sections are coded at different rates, depending on the channel quality, equipment and other factors.

The SIGNAL portion contains a PLCP header 211. The header includes length (LENGTH) 213 and data-rate (RATE) 215 information which can be used to calculate the total length of the frame in time. The final portion of the header 211 the SERVICE bits 221. Since this may be sent at a different data rate, it is part of the DATA section 219 of the frame.

The DATA section 219 in addition to the SERVICE bits includes a PSDU (PLCP Service Data Unit (SDU)) 223, a short set of tail bits 225 and pad bits 227 to isolate this frame from the next. . The PSDU payload 223, as shown in Figure 2B contains a "Duration" field 231 indicating the time remaining for this transaction to complete, measured from the end of the current frame. The PSDU also contains address fields 233, user data 235 and any of a variety of other desired fields, depending on the application. As shown the PSDU ends with the 802.11 FCS (Frame Check Sequence).

The PSDU can take a wide variety of different forms, the example of Figure 2A is provided as one possible example. The PSDU may include Logical Link Control (LLC)) authentication protocols, communications, operation, and management data and more. As can be seen in Figures 2A and 2B, 802.11 provides a standardized system for encapsulating user data of a variety of different types within 802.11 frames. The particular structure and format of the user data will depend upon the particular application. Internet Protocol, Ethernet, User Datagram Protocol and many other formats may be used in the SDU as allowed by the Specification. A variety of conventional frame formats allow different types of data, including telemetry data from the device to be transmitted over a variety of different network protocols. Such frames can be transmitted to and from the mobile device.

As can be seen in Figures 2A and 2B, there are two headers, one for the PLCP and another for the MAC. With other types of frames and packets the headers will differ and may have different names, such as preamble, information, configuration, etc. In the context of the present description, header is used to refer to all of the bits that describe the packet and its payload except for the actual payload. Header is used regardless of any particular name that might be given these fields in any particular protocol.

Figure 3A shows an example hardware configuration that may be used for the mobile device or a portion of the mobile device 101 of Figure 1. In this example, this portion of the device is able to function as a radio frequency identification tag. The tag has a controller, or microprocessor 301 to manage the transmit and receive operations described above. The controller is coupled to a modulator 303 to modulate any data that the controller is to send. The modulated data is sent to a power amplifier 305 that sends the modulated signal to an antenna 309. The amplifier controls the power used to send the signals and provides the various power levels mentioned above.

In addition to the transmit chain, the tag also has a receive chain that includes a low noise amplifier 311 coupled to the antenna to amplify any received signals that come through the antenna 309. The amplifier is coupled to a demodulator 315 that provides the demodulated data to the controller. The transmit and receive chains may also include additional components, such as oscillators, mixers, up and down converters and other components as may be desired for a particular radio frequency, modulation, and encoding scheme.

The receive chain also has an amplitude detector 313 to determine the received signal strength of the received signals. The determined values are used for the AP selection and location determination as described above. Alternatively, or in addition, other aspects of the signals may be measured and provided to the controller, these include carrier, interference, and noise levels, error rates, timing and other measures. The particular construction of the devices for making these measurements may be adapted to suit the particular radio environment and the types of signals to be measured.

The controller is coupled also to a transmit power controller 307, and the amplitude detector 313 and can additionally be coupled to any of the other components to receive or provide data and to provide control over the overall system. These connections are not shown. Instead, the connections are shown only to illustrate the path of incoming and outgoing messages. In the example of Figure 3A, the controller generates and interprets messages, counts the responses, counts the number of responses, makes comparisons and performs similar operations as described above.

The controller is further coupled to memory, the memory includes non-volatile memory, such as ROM (Read Only Memory) for program instruction and identification

values. The memory also includes writeable registers for storing measurements, operands, AP identifications and other values. The memory 317 shown in Figure 3A is coupled to the controller and includes both types. This memory can instead be included within the controller depending on the circumstances.

Figure 3A further shows a battery 319 for powering all of the components. The battery may be a conventional chemical cell, a photovoltaic cell, or any other type of power supply or combination of different types of power supplies. As mentioned above, in many applications, the tag is supplied with a single battery and it has a limited amount of total power available. Accordingly, by reducing the power consumed by the receiver and receiving functions of the tag as described above, the life of the battery can be extended.

Figure 3B shows additional details of the receive chain in the tag of Figure 3A. As in Figure 3A, a signal is received at an antenna 309. The antenna signal is applied to the LNA 311 and the output of the LNA is applied to variable gain stages (VG). In this example, there is an RF (radio frequency) variable gain amplifier (VGA) 323 coupled to a down converter or mixer 325 coupled to a baseband VGA 327. The baseband signal from the variable gain stage amplifier is applied to an ADC (analog to digital converter 329 and then to the demodulator 315 as shown in Figure 3A. A synthesizer 321 supplies a synthesized analog clock signal waveform output to the ADC 323 and the mixer 325. This particular receive chain is provided as an example configuration and the specific details may be adapted to suit any particular application.

As mentioned above, power savings are obtained by deactivating some or all of the receiver between frame preambles. In a simple design, all of the elements of Figure 3B can be turned off. Because some components take time to start back up and settle into a stable state, all of the components would then be switched on some time before they are needed. For example, the synthesizer in some embodiments, might take a millisecond to stabilize. The LNA and VGA might take tens of microseconds, and the ADCs fractions of a microsecond (hundreds of nanoseconds).

Moreover, if a low power timer or similar device is available for the system, the major part of the chip itself, incorporating the microprocessor plus media access control (MAC), can be shut down. These can be shut down, for example, by simply halting the

clock, or by even removing power to these components. Halting the clock can reduce digital power by as much as 20% for a very efficient microprocessor. The microprocessor can be restarted in a single clock cycle by restoring the clock. This time will be on the order of nanoseconds. For longer periods, the power to the digital core can be shut off. To activate the system, the system must then reboot the microprocessor and its operating system. This can take on the order of tens of milliseconds. This is generally much longer than any frame, but might be used if the frames are very far apart.

Accordingly, the receiver can be shut down in a variety of different ways. For very short intervals (fractions of a microsecond), the ADC can be deactivated between frame preambles. For longer intervals (microseconds), the ADC, the LNA and the VGAs can be deactivated. For a typical 802.11 frame, there is normally sufficient time between frames to deactivate, wait, and then reactivate these components. For even longer intervals (milliseconds), even the synthesizer can be deactivated. Since the PSDU of a frame can vary greatly in length a decision can be made, for example in the microprocessor or a separate power control processor how long the interval between frame headers will be and then the appropriate components can be deactivated to conserve power. In all of these scenarios the processor clock can be halted for additional power savings.

Figure 4 is a process flow diagram showing an example receiver activation and deactivation according to an embodiment of the invention. At block 401, a radio receiver is activated to receive frames on a wireless radio network. If the mobile device is already on the network, then this can occur based on a predicted time of arrival for a frame from a selected AP. If the mobile device is initiating a connection or has just moved to a new location, then this can be at a random time.

As in the other figures, while the transmitting station is characterized as an AP. This will depend on the particular application. In the standards for 802.11, a STA (station) is used. In other systems, the transmitters may be designated as base stations, routers, nodes, fixed terminals or by other terms. Embodiments of the present invention can provide benefits with a variety of different types of transmitters.

At block 403, the activated receiver listens for preambles of frames transmitted by radio transmitters. This can be either a particular AP, or any AP, depending on whether the receiver already has been configured for the particular wireless network.

At block 405, the receiver detects and receives a frame preamble. The preamble is then evaluated to determine if it is sent from an AP or other access node. If so, then it can be evaluated to determine the length of the frame. If it is not, for example, if it from another mobile device, then it is ignored. As mentioned above, the frame length can be determined based on information in the frame, such as the user data headers of Figures 2A and 2B. It can also be determined by measuring the length of other frames and using those measurements for an estimation of the length of all frames or some portion of all of the frames. Additional actions can also be taken such as measuring its signal strength, time of arrival, evaluating other header information etc.

At block 407, the receiver is deactivated. The entire receiver and the processor can be deactivated or, as mentioned above, only a portion of the receiver can be deactivated. The particular selection of components to activate and deactivate depends on the particular application and the design characteristics of the components. Faster component designs allow more components to be deactivated for shorter intervals. Another consideration is the importance of power savings as compared to quick responsiveness.

At block 409, the wait time to the next frame is determined. This may be done before the receiver is deactivated or afterwards, depending on which parts of the device are powered down when the receiver is deactivated. At block 411, the mobile device waits for the next frame, typically using a countdown timer. However, other types of timing devices may be used and at block 413, the receiver is again activated. The receiver can be activated enough in advance to be ready to receive the preamble of the next intended frame.

At block 415, the preamble of the next intended frame is received. In many applications, this will be the subsequent frame from the same AP. It may also be simply the next frame to occupy a channel, or the next frame from an AP in the system. The particular frame that is received will depend upon the determinations made when setting the wait timer. It should be noted that the receiver has been deactivated for some period

of time between block 407 and block 413 and this is a substantial part of how power consumption is reduced. As mentioned above, the duration of the wait is determined, at least in part, based on the information obtained in the header of the previous frame in block 405.

At block 415, the subsequent frame preamble is received by the reactivated receiver. At block 417, the address is read from the header to determine whether this subsequent frame is the intended frame and then at block 419, one or more quality parameters of the frame can be measured. If the address does not correspond to the desired frame, then the receiver can be deactivated until the next frame header can be received. An undesired frame may be one from a different AP or one from a different type of network node.

At block 421, the receiver is again deactivated to continue the power savings. The process can then end, return to the START, or return to block 411 to receive additional frames. The deactivation can happen as soon as the quality measure has been made and the address has been read.

For gaining network access or determining location, the address normally will be the from-address or some address that indicates the AP that transmitted the frame. After the mobile device has logged onto the network, the address may be the to-address, so that the mobile device is listening for and receiving only the frames that are intended for it. For other operations, other aspects of the frame header may be used, depending on the circumstances.

There are a variety of different quality parameters that may be measured depending on the application and the radio protocols and environments. In 802.11 systems, RSSI is used for many purposes. However, bit error rates, signal to noise, signal to interference, different amplitude measures and other measures of quality can be measured.

The information obtained by the mobile device can be used in a variety of different ways, depending on the application and the current status of the device. For example, the mobile device can accumulate RSSI measurements from the frame preambles and populate a table of AP addresses and RSSI. The table can be used to

select a network access node with which to register, for location determination or for a variety of other purposes.

Background Introduction

Addressing the lowest possible system power for Wi-Fi usage requires more than just a focus on active receive and transmit power. Some considered thought about the use cases of client Wi-Fi devices will lead to the conclusion that optimizing other factors, in addition to receive and transmit, allows for a system optimized design that can deliver average power usage below 10mW.

In order to deliver state of the art power handling, the system design needs to be attacked simultaneously on three levels: Circuit (Analog, Digital); 802.11 Mac/Phy (Fast Napping, Not Me); Systems/architectural (standby, fastboot, switching regulators)

Circuit Design

Circuit design on the analog and RF side is attacked with state of the art low power analog design techniques. An example of this from G2 is an ultra low power ADC at 44MHz that consumes 0.95mA. Its figure of merit is 160 pJ per conversion step designed across worst case process, voltage and temperature.

On the digital side, the main weapon is voltage scaling. The informed reader will know that at 90 and 65nm process nodes, voltages have ceased to scale linearly. The nominal supply voltage for a low leakage 65nm process is 1.2V - the same voltage as the generic 130nm node. Process scaling still gives a reduction due to capacitance reduction of the gate and routing capacitance but this only translates to about 20% power reduction per node.

Clock processing speed for 802.11g/a OFDM does not have to exceed 80MHz. Relative to the capability of the process node this is a low clock rate requirement. We take advantage of this to lower the supply voltage of the digital subsystem. A supply of 0.8V versus the nominal 1.2V is a design target.

If this voltage is varied and supplied via a linear regulator, the power reduction is proportional to voltage, i.e. the reduction would be 33%. However, if the voltage is varied and supplied through a switched regulator then the power reduction is proportional to voltage squared, i.e. the reduction would be 55%.

Since the analog subsystem will need a higher voltage to meet its dynamic range requirements, it is hard to justify separate system wide switching regulators for these multiple voltage requirements. However, when coupled with the architectural design that includes two on-chip switching buck regulators for the analog and digital systems, the G2 chip is capable of delivering the greater power reduction of 55% while maintaining a simplified system design.

The digital cell library design is also a factor in optimizing power consumption. Three special techniques can be used depending on design requirements.

First, G2 has a highly customized cell library used for the always-on section of the design. This library delivers the ultra low leakage of under $3\mu\text{A}$.

Second, for 65nm designs, leakage of the main digital cell library becomes a concern. Normally a library with higher than nominal V_t is used to ameliorate the problem. Lowering the supply voltage as discussed above also helps to lower the leakage. There is a tradeoff between the V_t used and the voltage that the chip can be operated at. The tradeoff is between leakage power and operating power. The optimal design point can be improved by the design of a customized cell library that allows operation at a lower supply voltage with the same V_t than might otherwise be possible.

Third, G2 has a flip flop design that allows for 'clock on demand' clocking, in chip modules with relatively low toggling rates this design significantly reduces the power contributed by the flip flops in the design. A novel feature of the same design is that it allows for sloppy clock delivery - reducing the risk of hold time failures - a problem made more likely when operating at ultra low supply voltages.

802.11 MAC/PHY

There are optimizations that can be made for power in the 802.11 MAC/PHY/RF system. Two examples of this are techniques G2 calls Fast Napping and Not Mine.

Fast Napping takes advantage of G2's expertise in analog/RF design that allows the analog system to take fast naps in the micro-second range. Doing this effectively allows a reduction of 50% of the analog, 100% of the PHY and 50% of the MAC. This allows for significant power reductions especially when the system is 'waiting to receive'.

Not Mine is a technique that includes hardware decoding of the 'to' address and 'duration' fields. When this data becomes available and is not for this device, it allows

rapid shutdown of the receiver until a predicted time in the future when the current packet transmission will be done and acknowledged. Power saving for this technique depends on the data encoding rate and the length of the data payload in the packet. For example when in the presence of other clients performing video streaming of 1512 byte packets, the power reduction can be enabled for 86% of the time.

Systems/Architectural

Techniques used by G2 in this category include a design architecture that allows low standby, fast booting and integration and control of switching regulators.

The combination of a very low standby and a fast boot mechanism allows the G2 Wi-Fi family to maximize the significant power advantages to be gained from the Wi-Fi power save poll mode and the newer techniques of the Wi-Fi Alliance's WMM-PS modes. This will be discussed in more depth below, suffice to say that Wi-Fi systems that cannot support low standby and fast-boot cannot take full advantage of the power saving potentially available.

Fast booting is enabled by several design features of the architecture, fast crystal startup, integrated control of the switching regulator bringup, booting of ecos in parallel to switching regulator settling, ROM coding of critical applications and optimizations to allow fast calibration setting of the RF subsystem.

G2 has integrated buck and boost switching regulators since its first production silicon. Integrating high power switching regulators, while maintaining a low noise environment for RF reception, is a significant design challenge that G2 has solved. Integrating the switching regulators means that the power quoted by G2 is 'as measured from the battery terminal'. This is appreciably different from chip vendors that quote power at the regulated supply pin of the chip, the system implications of supplying these voltages has significant power and cost penalties.

In addition, as mentioned above, the control of the buck regulators allows G2 to extract the maximum power saving from aggressive voltage scaling of the digital subsystem. Without an integrated power management architecture that includes the on-chip regulators this saving would not be possible in a cost effective manner.

States

There are 6 major operating states for a Wi-Fi chip, each of the above techniques can be applied to some but usually not all these states. The table below describes these states, the power expected for a 65nm G2 chip in each state and which techniques described above that are being applied.

State	G2 Chip Power	Techniques Applied
Sleeping	10uW	Ultra low leakage.
Waking & Booting	12mW	Fastboot
CPU Running (no radio)	21mW	Circuit and switching techniques.
Waiting to receive	40mW	Circuit, switching and fast napping techniques
Full Receiving	50mW	Circuit, switching and not mine techniques
Transmitting	60mW	Circuit Techniques. This is for a 0dBm transmit.

Table

A note on transmit: G2's current production silicon includes 14dBm PAs (Power Amplifiers). However a trend is for high efficiency III-V semiconductor external PAs to be deployed. These are very efficient relative to CMOS PAs. It is likely that G2's 65nm product will utilize a 0dBm PA driving an external high-efficiency PA in a module. It should be noted that in several cases transmitting is not the primary mode of operation - optimizing receive power generally delivers greater system benefits than optimizing transmit power.

Because each state can make use of different power techniques the overall power dissipated is dependent on the amount of time the chip spends in each state, i.e. on the operating scenario it is deployed in.

Operating Scenarios

There are many possible applications for a Wi-Fi chip but in general these applications slot into three operating scenarios, low-latency, high-latency and high-bandwidth. Examples of applications under each scenario are:

Low-Latency: voip, push to talk, browsing, Wi-Fi remotes, audio streaming

High-Latency: Wi-Fi idling, Wi-Fi sensors and actuators, wake-on-Wi-Fi, tags

High Bandwidth: receiving video stream, transmitting pictures from a camera

A Wi-Fi device can of course jump between these operating scenarios as required.

Low-Latency and High-Latency operating modes take advantage of legacy power save poll (PSP) or the newer WMM-PS modes of operation. The following will concentrate on legacy PSP modes but WMM-PS modes operate in similar fashion with similar benefits. See section 'Power Save Polling - background information' for background information on power save poll operation.

Low-Latency

In the low latency scenario the Wi-Fi device wakes for every DTIM (102ms if DTIM=TIM) to detect if the AP is holding traffic for it. If traffic is available the Wi-Fi device stays awake to receive the information, respond and go back to sleep. The steps the Wi-Fi device goes through are: Sleep; Wake; Wait for Beacon; Receive Beacon; Set timer and go to sleep.

If there is data for the device indicated by the TIM then after receiving the beacon the device will change modes to a high-bandwidth state, wait for data (broadcast or unicast), receive and ack data, set a timer, enter power save mode and go back to sleep.

Many Wi-Fi devices are not capable of being in a deep sleep state and waking rapidly. They simply cannot wake-up fast enough to respond in the 100ms time frames. These devices sometimes have a 'power save' mode from which they can wake quickly but this mode is relatively expensive due to high standby currents. The G2 silicon allows the chip to sit in a very low standby state, wake quickly and respond all within the 100ms time frame of a beacon period.

Let's take the case of a VoIP enabled Wi-Fi device. The device spends significant time waiting for a call, in this state the device wakes, finds no data for it and goes to sleep. The table below shows that the average power to achieve this with the targeted 65nm G2 Eden product is 3.34mW. In a cell phone with a 3.6V lithium ion rechargeable battery (710mA-hr), such as the Motorola Razr, this translates to 31 days of operation.

High-Latency

High latency applications allow applications such as home control. In this situation a similar technique to the low-latency scenario is used except that the chip does

not wake for every DTIM. Access Points will typically hold data for a power save client for up to 60 seconds. This means that the device can sleep through up to ~600 beacons. The table below shows operation at 50 Beacons (~5seconds), this would allow a home control device to operate from two Alkaline 1.5V AA cells for 8 years.

High Bandwidth

In this mode the chip is fully on and it comes down to the 50mWs of receive power delivered with advanced circuit and switch regulator techniques described above. It should be mentioned however that a Wi-Fi chip in a client device will not spend all its life in this mode, on average there are still very long periods when the device is idle and can be placed in the low-latency mode.

Power Save Polling - background information

Infrastructure Network power management

Power saving is achieved by having the AP buffer frames for an associated station while the station is asleep. Stations can indicate to APs that they are using power management and APs can indicate to stations the status of their buffer.

A station can be in either Continuously Active Mode (CAM) or Power Save Mode (PSM). A station notifies its mode to an AP using the Power Management bit in the Frame Control field in transmitted frames. Setting the bit indicates PSM and clearing the bit indicates CAM.

APs indicate the buffer status in Beacon frames which contain a Traffic Indication Map (TIM) information element. There are two types of TIM, a TIM and a Delivery Traffic Indication Map (DTIM).

There are two different cases of frame buffering, these being for unicast and multicast frames.

Unicast frame buffering and delivery

When a station associates with an AP, it indicates a Listen Interval in the Association Request frame. The Listen Interval tells the AP how frequently the station intends to listen to beacon frames for indication of buffered unicast frames. The Listen Interval is a 16 bit value and is specified in multiples of the beacon period. This value gives the AP a clue to the buffer requirements of the station and it is conceivable that an

AP could reject an association if the Listen Interval is too long. An AP is entitled to discard frames that are buffered for greater than the Listen Interval.

The Association Response frame sent from the AP to the station contains an association identifier (AID). The AID is unique for each station associated with the AP and is used by the AP to indicate when buffered unicast frames are present. The AP maintains a Virtual Bitmap of 2008 bits where each bit represents the buffer status for an AID. The Partial Virtual Bitmap in the TIM contains one or more octets from the Virtual Bitmap representing those stations for which it has frames buffered. For example if the AP has frames buffered for AID 1, bit 1 of the first octet will be set to 1. Bits 1 to 7 in the Bitmap Control octet are the Bitmap Offset; this is used to allow stations to determine which parts of the Virtual Bitmap have been included in the TIM.

PS-Poll

To retrieve buffered frames, a station has to wake up periodically and listen to beacon frames. If it determines that the AP has buffered frames for it, the station sends a PS-Poll frame to the AP. In response to the PS-Poll frame, the AP will send a single frame from its buffer. The AP should not remove the frame from its buffer until it has received an ACK for the frame. The AP indicates if more frames are buffered using the More Data bit in the Frame control field; there are more frames if the bit is set. The station continues to send PS-Poll frames to the AP until either the More Data bit is cleared in the received frame or the bit corresponding to its AID in the TIM is cleared. The station should empty the AP buffer before going back to sleep.

PSM/CAM

There is nothing to stop a station switching between PSM and CAM modes. Another method for a station to retrieve buffered frames is to switch from PSM to CAM when it wakes up. The AP will then send any buffered frames to the station without requiring the station to poll for each frame.

In this mode of operation, the AP isn't required to set the More Data bit, so the station needs a mechanism to determine when it can go back to sleep. One method is to use a timer to detect a receive idle period. The timer is restarted whenever a frame is received and if the timer expires, the station assumes that no more frames are going to be received. After receiving all its frames, the station can return to PSM, send any frames it

has been waiting to send and then go back to sleep. Note that when going back to PSM mode, every frame sent needs to have the Power Management bit set otherwise the station will be toggling from PSM to CAM.

Broadcast and Multicast frame buffering and delivery

Broadcast and multicast frames are also buffered by the AP for stations in PSM. The broadcast and multicast frames are sent immediately after a beacon containing a DTIM. Stations wishing to receive these frames need to be awake when the DTIM is transmitted so that they can determine if any frames are about to be delivered.

The AP sends a DTIM at a rate determined by the DTIM period. The DTIM period is indicated in the TIM DTIM Period field, this is a single octet, so the DTIM period can be between 1 and 255. The DTIM period of 0 is reserved. The DTIM period is specified in units of the beacon interval, the beacon interval is commonly set to 100 ms however it is configurable. The DTIM period is typically set between 1 and 3.

Beacon frames additionally contain the DTIM count in the TIM. The DTIM count indicates how many more beacons, including the current one, appear before the next DTIM.

The AP indicates that it has broadcast or multicast frames for delivery by setting the Traffic Indicator bit, bit 0 in the TIM Bitmap Control field. This bit is associated with AID 0.

A station in PSM should wake to receive the DTIM, check the Traffic Indicator bit and receive buffered frames if the bit was set, otherwise it may go to sleep. APs don't set the More Data bit when sending buffered broadcast or multicast frames, so the station needs a method to determine when to return to sleep. A receive idle timer could be used as described in the PSM/CAM section above.

Napping of receiver to save power

The receiver can be operated intermittently with a very short cycle time (eg. 500ns on, 2µs off; always shorter than the signal preamble so you don't miss a signal) until a signal is received. ("Napping" refers to the time the receiver is asleep e.g. for 2µs.)

Extensions:

(a) Retaining power to the blocks such as DC offset correction or synth. Loop filter that have long settling times while other blocks are switched off.

(b) Detecting a signal by the received power level increasing since the last reception (eg. 2.5 μ s ago), rather than it exceeding a threshold. This exploits the facts that (i) we are only interested in a new signal preamble, (ii) 802.11 needs SNR > 0dB to receive a signal so a new signal must raise the power level, and (iii) 802.11 signals have fairly "square" power versus time plots so the start of a new signal is fairly easy to detect.

(c) Using dedicated RSSI hardware (in our case, three power comparators) to allow AGC to take place more quickly, freeing up more time for intermittent operation without missing a frame. This is almost essential given the very short (8 μ s) 802.11 short training sequence.

WiFi MAC layer ex-protocol power saving mechanism by sleeping the PHY when medium unavailable or of no interest.

Background:

WiFi (802.11) receivers typically demodulate all frames which they receive. This is done both while waiting for a frame for the particular station (STA), and while waiting for an opportunity to transmit. Both waiting (without demodulating) and demodulating incur a cost in energy dissipation. In a case where the STA is spending a significant amount of its time listening for frames or waiting to transmit as compared to transmitting, this energy can become a dominant contributor to system energy dissipation.

The PHY can be put to sleep for cases where 1) The medium has a transaction of no interest to us, or 2) The medium is quiet, or 3) There has been a collision on the medium.

Example Implementation

WiFi (802.11) frames contain information on the type and length of the current frame and the current transaction, and a MAC address identifying the recipient of the current transaction. Both of these units of information are stored near the start (in time) of the frame as transmitted over the air. By examining these units of information, an intelligent STA can determine whether the current transaction is of interest, as well the length of time for which the current transaction will be using the medium. If the current

transaction is of no interest to the STA (principally that it is either a direct packet not addressed to the STA, or a broadcast packet of no interest to the STA), then the STA has enough information to turn off the PHY (baseband/analog/RF) until such time as the medium will again be available (either for transmitting a frame or for the possible reception of a packet of interest).

The information which allows a STA to do this is:

- 1) length and rate of current frame
- 2) DURATION field of the frame
- 3) TYPE field
- 4) ADDR-1 field

In addition, to take advantage of this a STA must also be able to calculate the length in μs of the current frame, as well as identifying the time at which the current frame went onto the air. These are commonly needed anyway.

This method will apply both to direct and broadcast data frames, as well as to RTS-CTS frames. The benefits will be greatest in an environment where long frames are being transmitted, as this provides the longest amount of time for the PHY to sleep relative to its awake time. Of course, in an environment where there is very little traffic (say no traffic other than for the STA), then there is little benefit from sleeping on "not for us". In this case an alternative method for sleeping becomes viable, as discussed in extensions ("sleep while quiet"). In a very busy environment with many collisions, this method will fail to provide the required information to sleep, in which case a further extension ("sleep on collision") can be used.

Extensions:

"Sleep while Quiet"

In the case where the medium is very quiet (i.e. little traffic) the STA (while waiting for a frame) could duty cycle the receiver on a millisecond or greater time scale. The exact time scale would be chosen based on knowledge of the frame lengths it has seen or expects to see, so that the likelihood of missing a frame (keeping in mind that there will be retries for missed frames) is reduced to a desired likelihood. This duty cycling of the baseband/analog/RF PHY could provide significant savings. The duty cycle could be chosen so that the PHY listens for a period larger than the SIFS time or

DIF times or even minimum backoff time (in a quiet medium the Contention Window will be small) to increase likelihood of seeing a packet if it is there.

Successful reception of a packet is not necessary to indicate the presence of activity on the medium - even detection of a frame mid-way through provides sufficient information to change the duty cycling ratio to a higher value so that the next frame will be received from the start (which is to say, successfully) - in fact after seeing the energy, the STA may simply go back to its normal mode of operation until it either receives a frame, or has waited long enough to determine that the medium is still quiet and return to its "sleep while quiet" mode.

"Sleep in Collision"

In the case of a collision between two frames on the medium, the 802.11 protocol indicates that stations must back off for an Errored Inter Frame Space. We can extend our invention so that in the case where the STA sees significant energy on the medium which it cannot reliably demodulate (e.g., fails to find valid preamble etc), then we can assume that there has been a collision, and put the PHY to sleep either for EIFS, or for the estimated length of time for which the medium will be busy. This estimated time can be obtained from keeping track of the mean length (in μs) of frames on the air, and using that time or some fraction of it.

In the event that the PHY sleeps for longer than was needed and a frame is missed, there is a small but not distressing penalty in energy - the transmitter of the missed frame will attempt to transmit it again - the penalty is that we need to wait longer for that frame. This too can be tracked by the STA to tune the length for which it stays asleep after a collision. Retry of a frame transit is indicated by setting a flag in the transmitted frame - if the STA sees that it is receiving "retries" consistently after sleeping due to a collision, it can reduce the time for which it sleeps after a collision until it is no longer seeing a significant number of retries after collision relative to the normal retry rate it sees. This retry rate must also be compensated for length of frame and rate - both of which affect reliability of reception.

Unique Architecture

The existence of all necessary functions on a single integrated circuit including higher layer software such as tcp/ip and security supplicants and including the power

management with on-chip linear and switching regulators and the above circuit and protocol techniques allows the chip to respond to events at speeds not possible in multi-chip systems where the higher layer software control is conducted in a host processor in the system. This means that the power in the overall system can be reduced because the host processor can be removed or at least reduced in cost, complexity and power needs.

An example of activity possible in a fully integrated chip not possible in a multi-chip system is the ability of the WLAN system to operate autonomously in and make decisions such as when and how to roam from one radio access point to another. Because everything is integrated and tightly controlled, this roaming can be achieved without loss of streaming media data as experienced by the user of the device.

The approaches described herein may be used in any situation in which a mobile node is to select from among multiple other nodes, whether fixed or mobile, for wireless communication. For example, the approaches may be used when a single node is introduced into a new environment, when nodes, whether fixed or mobile, are added to or removed from an existing environment and when multiple mobile nodes seek to establish ad-hoc networks.

A lesser or more equipped transmitter or receiver than the examples described above may be preferred for certain implementations. Therefore, the configuration of the exemplary tag in Figures 3A and 3B or the environment in Figure 1 will vary from implementation to implementation depending upon numerous factors, such as price constraints, performance requirements, technological improvements, or other circumstances. The particular nature of any attached devices may be adapted to the intended use of the device. Any one or more of the devices, interfaces, or interconnects may be eliminated from this system and others may be added. For example, a variety of different connections to the access point may be provided based on different wired or wireless protocols.

In the description above, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form.

The present invention may include various steps. The steps of the present invention may be performed by hardware components, such as those shown in the Figures, or may be embodied in machine-executable instructions, which may be used to cause general-purpose or special-purpose processor or logic circuits programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware and software.

The present invention may be provided as a computer program product which may include a machine-readable medium having stored thereon instructions which may be used to program an agent or a computer system to perform a process according to the present invention. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, magnet or optical cards, flash memory, or other type of machine-readable media suitable for storing electronic instructions. In the example of Figure 3A, a separate memory is provided. However, the memory may also be resident within the microprocessor. Moreover, the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer to a requesting computer by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

Many of the methods and apparatus are described in their most basic form but steps may be added to or deleted from any of the methods and components may be added or subtracted from any of the described apparatus without departing from the basic scope of the present invention. It will be apparent to those skilled in the art that many further modifications and adaptations may be made. The particular embodiments are not provided to limit the invention but to illustrate it. The scope of the present invention is not to be determined by the specific examples provided above but only by the claims below.

CLAIMS

1. A method comprising:
activating a radio receiver of a mobile device;
listening for preambles of frames transmitted by a radio transmitter;
receiving a frame preamble;
deactivating the receiver after receiving a frame preamble;
reactivating the receiver after a time determined based on the received frame preamble to receive a subsequent frame preamble;
receiving the subsequent frame preamble;
determining an address for the subsequent frame; and
measuring a quality parameter of the subsequent frame, if the address corresponds to the radio transmitter.
2. The method of Claim 1, wherein receiving the frame preamble comprises receiving a from address, the method further comprising determining whether the from address corresponds to a desired transmitter.
3. The method of any one or more of claims 1 or 2, wherein reactivating the receiver comprises reactivating the receiver after a time determined by estimating the duration between frame transmissions from the radio transmitter based on measurements of durations between previous frame transmissions.
4. The method of any one or more of claims 1, 2, or 3, further comprising storing the measured quality parameter in association with an identification of the radio transmitter.
5. The method of Claim 4, wherein the measured quality parameter is a received signal strength indicator.
6. The method of Claim 4, wherein the identification of the radio transmitter is a function of the received from address.
7. The method of any one or more of claims Claim 1.2, or 3, further comprising:
measuring the received signal strength of a plurality of frame received from a plurality of radio transmitters; and
selecting a radio for reception based on the measured received signal strength,
and

wherein receiving a frame preamble comprises receiving a frame preamble from the selected radio.

8. The method of Claim 7, wherein selecting a radio comprises comparing the measured received signal strengths and selecting the radio with the highest received signal strengths.

9. The method of any one or more of the above claims, further comprising measuring times of arrival of received frame preambles and determining a location relative to external radio transmitters based on the measured times of arrival.

10. The method of Claim 1, further comprising, after receiving a frame preamble, receiving the frame corresponding to the preamble and evaluating the frame and wherein deactivating the receiver comprises deactivating the receiver after receiving the frame.

11. The method of any one or more of the above claims, wherein deactivating the receiver comprises deactivating a receive amplifier of the mobile device.

12. The method of Claim 11, wherein deactivating the receiver further comprises deactivating a synthesizer, a downconverter, and a processing core of the mobile device.

13. A machine-readable medium containing instructions which when executed by the machine cause the machine to perform the operations of any one or more of the above claims.

14. An apparatus comprising means for performing each of the operations of any one or more of claims 1-13.


15. An apparatus comprising:

a receiver to listen for preambles of frames transmitted by a radio transmitter, to receive a frame preamble, to receive a subsequent frame preamble, to determine an address for the subsequent frame, and to measure a quality parameter of the subsequent frame, if the address corresponds to the radio transmitter;

a processor to deactivate the receiver after receiving the frame preamble, and to reactivate the receiver after a determined time based on the received frame preamble to receive the subsequent frame preamble; and

a timer to count the determined time.

1/6

FIG. 1




FIG. 2A

223 ↘





FIG. 2B

FIG. 3A




FIG. 3B

6/6

FIG. 4