



(51) International Patent Classification:
H04L 7/02 (2006.01) *H04L 7/033* (2006.01)
G06F 1/32 (2006.01)

(21) International Application Number: PCT/EP2009/067814

(22) **International Filing Date:** 22 December 2009 (22.12.2009)

(25) **Filing Language:** English

(26) **Publication Language:** English

(30) Priority Data:
08254192.1 30 December 2008 (30.12.2008) EP

(71) Applicant (for all designated States except US): **ST ERICSSON SA** [CH/CH]; 39 Chemin du Champ-des Filles, 1228 Plan - des Ouates, Geneva (CH).

(72) Inventors; and

(75) **Inventors/Applicants (for US only): PELTONEN, Janne** [FI/FI]; Tiilcarinkatu 12, FI-24240 Salo (FI). **RAUHALA, Antti** [FI/FI]; Pajusirkuntie 2, FI-20760 Piiroinen (FI).

ispanristi (FI). **SALMI, Mika** [FI/FI]; Ylannekatu 5A,
FI-20540 Turku (FI).

(74) **Agent: STYLE, Kelda Camilla Karen**; Page White & Farrer, Bedford House, John Street, London Greater London WC1N 2BF (GB).

(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, CL, 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, PE, 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.

(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,

[Continued on next page]

(54) Title: CLOCKING SIGNAL CONTROL




Figure 3

(57) Abstract: A system (21) comprises a first circuit (23) for transmitting information to a second circuit (22). The first circuit (23) comprises a signal generator (34) arranged to generate a signal (31) for transmission to the second circuit (22). The signal (31) has a first state for causing the second circuit (22) to perform a first operation and a second state for causing the second circuit (22) to perform a second operation. The signal generator (34) comprises means to modulate the generated signal (31) during a data burst period (52) to encode data into the generated signal (31). The data burst period (52) occurs within a period in which the generated signal (31) is in the second state. The signal (31) is modulated such that the encoded data is distinguishable from the generated signal (31). In an embodiment, a data burst period (52a, 52b) comprises a plurality of data periods (54) during which the state of the signal is representative of the data encoded in the signal (31). The data periods (54) may be separated by padding periods (55) during which the signal (31) is in the second state. The data burst period may be initiated following a delay after the signal changes from being in the first state to being in the second state.



TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, **Published:**

ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM,

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,

ML, MR, NE, SN, TD, TG).

— *with international search report (Art. 21(3))*

Clocking Signal Control

The present invention relates to the communication of information between sub-blocks of a system, for example information relating to the control of
5 clocking signals in a wireless cellular system.

Many systems, such as wireless cellular systems, typically comprise several logical and physical sub-blocks. These sub-blocks can be, for example, ASIC (Application Specific Integrated Circuit) units forming the total working system.
10 Each ASIC performs a specific function within the system and provides the best and most natural environment for the set of related logical operations required to perform that function. For example, the functions relating to power management (PM) may be provided by one ASIC while the functions relating to RF (Radio-Frequency) processing may be provided in another separate
15 ASIC. Due to the modular nature of this type of arrangement, it is necessary that information can be effectively communicated between the ASIC units. Information is typically transmitted between ASICs on digital lines which terminate at the ASICs by means of electrical connection pins or pads. Typically, the number of pins or pads in an ASIC is limited and each pin is pre-
20 allocated to a specific function.

One kind of information that needs to be communicated between ASICs relates to the control of clocking signals. One of the typical functions of a PM ASIC is to control entry into and exit from 'sleep mode' in which clocking
25 signals normally provided to the various parts of the system are temporarily switched off in order to conserve power. However, in wireless cellular systems, clock generation and clocking functions (including operations to actually perform entry into, and exit from, sleep mode) are typically performed inside the RF ASIC. For this reason, in order for the PM ASIC to control entry
30 into and exit from sleep mode, by instructing the RF ASIC to do so, it is necessary for the PM ASIC to communicate with the RF ASIC.

In order to instruct the RM ASIC to enter and exit sleep mode, the PM ASIC typically transmits a two-state signal on a single digital line to the RF ASIC,

which has a pin dedicated to receiving the signal. The state of the signal depends on the detection by the PM ASIC of one or more clock request signals which are asserted by various components of the system when clocking signals are required. If no clock request signals are asserted then the system is considered to be inactive in which case the sleep mode is entered to conserve power. The signal transmitted from the PM ASIC to the RM ASIC is placed in a first (e.g. low) state causing the RF ASIC to not provide clocking signals to the various components of the system. When at least one clock request signal is subsequently asserted by at least one of the components, sleep mode is exited from. The transmitted signal is placed in a second (e.g. high) state causing the RF ASIC to provide clocking signals to each of the components.

As systems are becoming more complex, and in order to form adaptive and dynamically optimised systems, it is often necessary to communicate additional information between ASICs. However, it is usually not possible to add pins to accommodate the additional digital lines required to transmit this information due to limited physical size of the ASICs.

In accordance with one aspect of the invention, there is provided a first circuit for transmitting information to a second circuit, the first circuit comprising a signal generator arranged to generate a signal for transmission to the second circuit, the signal having a first state for causing the second circuit to perform a first operation and having a second state for causing the second circuit to perform a second operation; wherein the signal generator comprises means to modulate the generated signal during a data burst period to encode data into the generated signal, the data burst period occurring within a period in which the generated signal is in the second state, the signal being modulated such that the encoded data is distinguishable from the generated signal.

Embodiments of the invention will now be described, by way of example only, with reference to the Figures, in which:

Figure 1 is a schematic diagram of a system in which a clock enable signal is communicated between a first ASIC and a second ASIC;

Figure 2 illustrates the clock enable signal used in the system shown in Figure 1;

- 5 Figure 3 is a schematic diagram of a system in which a System Clock Request signal, SysClkReq, is communicated between a first ASIC and a second ASIC;

Figure 4a illustrates the SysClkReq signal used in the system shown in Figure 3; and

- 10 Figure 4b shows a close-up view of the data burst period of the signal shown in Figure 4a.

Figure 1 illustrates a system 1 in which a clock enable signal (illustrated in Figure 2) is communicated between a first ASIC 2 and a second ASIC 3 in order to control clocking functions, such as controlling entry into, and exit from, sleep mode. Figure 3 illustrates a system in which a System Clock Request signal, SysClkReq, (illustrated in Figures 4a and 4b) is communicated between a first ASIC 22 and a second ASIC 23. The SysClkReq signal has a form such that it can function as a clock enable signal as used in the system shown in Figure 1. However, the SysClkReq signal also comprises a message burst allowing it to convey additional information between the ASICs, such as specific clocking information.

Figure 1 schematically illustrates the system 1 comprising a first ASIC 2 and a second ASIC 3. In this example, the first ASIC 2 is a RF ASIC which performs clock generation and clocking functions, while the second ASIC 3 is a PM ASIC providing power management functions. One of the functions of the first ASIC 2 is to provide clocking signals 4a-c to one or more further ASICs 5a-c in the system 1. As shown in Figure 1, a separate clocking signal 4a-c is provided for each of the respective further ASICs 5a-c. The first ASIC 2 comprises an oscillator 6 which generates a clock signal 7 which is split into multiple signals from which the clocking signals 4a-c provided to the further ASICs 5a-c are derived. In the example shown in Figure 1, the further ASICs 5a-c are arranged in parallel. Each of the split clock signals 7 passes through

a respective gate 8a-c which can selectively allow or block the signal 7 depending on respective gate control signals 9a-c. The respective gate control signals 9a-c are generated by control logic 10 thereby allowing the control logic 10, and thus the first ASIC 2, to selectively provide the respective
5 clocking signals 4a-c to respective further ASICs 5a-c, or not.

The provision of clocking signals 4a-c to the further ASICs 5a-c requires significant power consumption. Therefore, if the system 1 becomes temporarily inactive so that none of the further ASICs 5a-c require clocking
10 signals 4a-c, in order to conserve power, the system 1 may enter a state in which no clocking signals 4a-c are provided to any of the further ASICs 5a-c. This mode of operation is often referred to as a sleep mode. A similar low-power mode may be entered in other circumstances, for example when the power supply is running low. It can be seen that the control logic 10 may
15 cause the system 1 to enter into, or exit from, sleep mode by "shutting" or "opening" the gates 8a-c using appropriate control signals 9a-c.

From the foregoing, it can be seen that it is a role of the first (RF) ASIC 2 to perform actual entry into and exit from sleep mode. However, since entry into
20 sleep mode is generally a power-saving mechanism, the actual control and responsibility of whether and when to enter into, or exit from, sleep mode is typically allocated to the second (PM) ASIC 3. This means that information needs to be transmitted from the second ASIC 3 to the first ASIC 2. Specifically, in order to instruct the first ASIC 2 to enter into, or exit from, sleep
25 mode, the second ASIC 3 outputs a signal 11, which is received by the first ASIC 2, and whose state defines whether sleep mode should be entered into or exited from. This signal is illustrated in Figure 2. For example, when the signal 11 is not asserted, or is in a low state, the first ASIC 2 causes the system 1 to be placed or maintained in sleep mode, whereas when the signal
30 11 is asserted, or is in a high state, the first ASIC 2 causes the system 1 to be placed or maintained in the normal, non-sleep, mode. The first ASIC 2 comprises a contact pin 12 which is dedicated to receiving the signal 11.

In the following description, the signal 11, and any equivalent signal, which controls whether or not clocking signals 4 are provided to parts of the system 1, are referred to as clock enable signals. The clock enable signal 11 has two possible states. The first state indicates that no clocking signals 4a-c derived from a system clock, such as oscillator 6, are to be provided to any part of the system 1 (sleep mode). The second state indicates that clocking signals 4a-c are to be provided.

As mentioned above, when the further ASICs 5a-c are inactive, the system 1 may be placed in sleep mode during which time no clocking signals 4a-c are provided to the further ASICs 5a-c. However, one or more of the further ASICs 5a-c may subsequently become active and therefore require clocking signals 4a-c once again. To this end, when any of the further ASICs 5a-c require clocking signals 4a-c, the ASIC 5 requiring the clocking signals 4 asserts a respective clock request signal 13. Any clock request signals 13a-c asserted by the respective further ASICs 5a-c are received by the second ASIC 3. The second ASIC 3 is arranged to produce the clock enable signal 11 according to the respective received clock request signals 13a-c. Specifically, if the second ASIC 3 receives one or more asserted clock request signals 13a-c, the clock enable signal 11 is placed or maintained in a high state, while if no clock request signals 13a-c are asserted, the clock enable signal 11 is placed or maintained in a low state.

In the system illustrated in Figure 1, the second ASIC 3 comprises logic 14, in the form of an OR gate, which receives each of the clock request signals 13a-c as inputs and produces an output 15 from which the clock enable signal 11 is derived. Figure 2 illustrates the form of the clock enable signal. The output 15 of logic 14, and thus the clock enable signal 11, is high if any of the inputs are high and is low if none of the inputs are high. It can thus be seen that the clock enable signal 11 remains high while any one of the further ASICs 5a-c is asserting a clock request signal 13a-c. This is the case during the period P₂ illustrated in Figure 2. During this period, the control logic 10 of the first ASIC 2 receiving the clock enable signal 11 causes clocking signals 4a-c to be provided to respective further ASICs 5a-c. Conversely, the clock enable signal

11 output by logic 14 remains low while none of the further ASICs 5a-c is asserting a clock request signal 13a-c. This is the case during the periods P_1 and P_3 illustrated in Figure 2. During these periods, the control logic 10 causes no clocking signals 4a-c to be provided to any of the further ASICs 5a-c.

In one arrangement, the control logic 10 is arranged to react to changes in the state of the clock enable signal 11 only after a delay. For example, the control logic 10 may be arranged to initiate entry into sleep mode when the clock enable signal 11 changes from high to low only after a certain delay period and only if the clock request signal 11 has remained low for the entire delay period. This condition may be imposed to ensure that sleep mode is not entered when the system is inactive for only a short time. In such cases it may be unnecessary to enter sleep mode, and due to the time required to perform the sleep mode entry and exit procedures (which may be longer than the period of inactivity), may actually result in significant system delays. In another example, the control logic may be arranged to initiate exit from sleep mode when the clock enable signal changes from low to high only after a certain delay period. This condition may be imposed to provide sufficient time to carry out any relevant wake-up initialisation procedures before clocking signals are provided. Since the provision of clocking signals consumes power, delaying this step until the initialisation procedure is complete conserves power.

Using the system shown in Figure 1, it is not possible for the first ASIC 2 to know which of the further ASICs 5a-c caused the sleep mode to be exited by asserting a clock request signal 13a-c. This is because the clock request signals 13a-c are combined by the logic 14 to derive the clock enable signal 11 in such a way that information regarding the origin of the clock request signals 13a-c is lost. The simple two-state clock enable signal 11 described above is insufficient to convey any additional information.

As new systems are becoming increasingly complex, some level of dynamical adaptation is desirable in order to optimise the system. In the system described above, it would be advantageous for the first ASIC 2 to know which

of the further ASICs 5a-c requested clocking signals. One simple way to achieve this would be to provide additional parallel digital control lines between the first ASIC 2 and the second ASIC 3 to convey the information. Another way would be for the clock request signals 13a-c to be provided
5 directly to the first ASIC 2. Both of these techniques require additional pins to be provided on the first ASIC 2 in order to accommodate the additional digital lines. As mentioned above, all available physical space for connection pins is typically already used for existing functions and so it may not be possible in some, but not all, situations to add pins. Furthermore, the amount of
10 information it is desirable to transmit between ASICs is becoming larger. For example, the number of components which may independently request clocking signals is becoming larger. A large number of additional pins would therefore be required to accommodate the digital lines necessary to transmit the clocking information.

15 Figure 3 illustrates a system 21 in which a SysClkReq signal 31 is transmitted to a first ASIC 22 from a second ASIC 23. As described in greater detail below, the SysClkReq signal 31, illustrated in Figures 4a and 4b, can be used in the same way as the clock enable signal 11 used in the system 1 illustrated
20 in Figure 1. However, the SysClkReq signal 31 also comprises a data burst 52 allowing it to convey additional information between the second ASIC 23 and the first ASIC 22. For example, this information may identify which further ASICs 25 in the system 21 have asserted a clock request signal 33. Since the SysClkReq signal 31 may be transmitted on a single line, no additional pins are required to transmit the additional information, or, depending on the
25 amount of information required, the number of additional pins required may be minimised.

The system illustrated in Figure 3 is similar to the system shown in Figure 1
30 except in the way that the SysClkReq signal 31 is generated by logic 34 and used by control logic 30. The system 21 shown in Figure 3 comprises a first ASIC 22, which in this example is a RF ASIC, and a second ASIC 23, which in this example is a PM ASIC. The first ASIC 22 outputs one or more clocking signals 24a-c which are provided to one or more respective further ASICS

25a-c. The further ASICs 25a-c are illustrated as being arranged in parallel although they could be arranged in any other suitable configuration. The first ASIC 22 comprises an oscillator 26 which generates a clock signal 27. The clock signal 27 is divided into separate signals which form the clocking signals 24a-c output by the first ASIC 22. The separate clock signals each pass through a respective gate 28a-c which block or allow the passage of the respective clock signals 7 according to respective control signals 29a-c output by the control logic 30. The control logic 30 operates in accordance with a SysClkReq signal 31 output by the second ASIC 23 and received by the first ASIC 22. The first ASIC 22 comprises a pin 32 for receiving the SysClkReq signal 31.

Each of the further ASICs 25a-c can request respective clocking signals 24a-c by outputting a respective clock request signal 33a-c. The second ASIC 23 comprises logic 34 which receives clock request signals 33a-c and generates an output 35 forming the SysClkReq signal 31.

The SysClkReq signal 31 will now be described with reference to Figures 4a and 4b. Figure 4a schematically illustrates the form of the SysClkReq signal 31 in one exemplary embodiment. The SysClkReq signal 31 shown in Figure 4a is similar to the clock enable signal 11 described above in relation to the system shown in Figure 1, except that the SysClkReq signal 31 comprises a data burst period 52 in the middle part of the signal. Figure 4b shows a close-up view of the data burst period 52 of the SysClkReq signal 31. During the data burst period 52, the clock enable component of the signal 31 is modulated to encode data into the signal 31. The signal 31 is created in such a way that the clock enable component of the signal can be distinguished from the serial data component by the control logic 30 of the first ASIC 22. This in turn enables the SysClkReq signal to function as both a clock enable signal and a data carrier.

At a time when no clock request signals 33a-c are being asserted, the SysClkReq signal 52 is generated in a low state, shown as period 50 on the left hand side of Figure 4a. When the first ASIC 22 receives the SysClkReq

signal 31 in a low state, the control logic 30 of the first ASIC 22 controls the gates 28a-c such that no clocking signals 24a-c are provided to the further ASICs 25a-c. In this state, the system 21 is in sleep mode.

5 At a time T_0 , at least one of the further ASICs 25a-c requests clocking signals 24 by asserting a respective clock request signal 33. The logic 34 detects the assertion of one or more clock request signals 33a-c and causes the SysClkReq signal 31 to be generated in a high state, shown as period 51 in Figure 4a. When the first ASIC 22 receives the SysClkReq signal 31 in a high state,
10 state, the control logic 30 of the first ASIC 22 initiates a system wake-up procedure to exit sleep mode including controlling the gates 28a-c such that clocking signals 24 are provided to one or more of the further ASICs 25. The first rising edge of the SysClkReq signal 31 at time T_0 signals that the wake-up procedure should be initiated.

15 In one embodiment, the control logic 30 of the first ASIC 22 is arranged so that the SysClkReq signal 31 must remain in a high state for a certain time period before the wake-up procedure is initiated and one or more clocking signals 24 are provided to one or more respective further ASICs 25. As shown
20 in Figure 4a, the SysClkReq signal 31 remains in a high state for a period ΔT , until a time T_1 , at which time the message burst portion of the SysClkReq signal 31 begins. A delay between the first rise in the SysClkReq signal 31 from a low state to a high state to the beginning of the data burst may be used for several reasons. First, the delay ensures that the SysClkReq signal 31 is
25 maintained in a high state for the time period required by the control logic 30 to initiate the wake-up procedure. Second, as described in greater detail below, the delay helps the control logic 30 to filter out, distinguish or otherwise separate the data component of the SysClkReq signal 31 for the purpose of using the SysClkReq signal 31 as a clock enable signal.

30 At time T_1 , a first data burst period 52a begins. During this period, data bits are encoded into the SysClkReq signal 31 allowing the SysClkReq signal 31 to be used to communicate data from the second ASIC 23 to the first ASIC 22.

In particular, as best seen in Figure 4b, the data burst period 52 comprises a series of data frames 53a-z, each data frame 53a-z conveying one bit of data. Each data frame 53a-z comprises a data period 54 in which the actual data bit is encoded, and a padding period 55 which separates consecutive data bits.

- 5 The padding periods 55 are indicated in Figure 4b in grey for ease of visualisation.

During the data period 54 of a data frame 53 the SysClkReq signal 31 is set to either a high or low state depending on whether the data bit to be encoded in that data frame is a logical 1 or logical 0. During the padding period 55 of the data frame 53 the SysClkReq signal 31 is set to or maintained in a high state. For example, the data frames 53b, 53c and 53d shown in Figure 4b encode binary values 1, 0 and 0 respectively since the signal is respectively in a high, low and low state during the respective data periods 54b, 54c, 54d. The signal is in a high state during each of the respective padding periods 55b, 55c and 55d. The padding periods 55 occurring between consecutive data bits makes it possible to filter, distinguish or otherwise separate the serial data conveyed by the SysClkReq signal 31 from the clock enable component of the signal 31. As described in greater detail below, the presence of the padding periods helps the control logic 30 to filter out the data component of the SysClkReq signal 31 for the purpose of using the SysClkReq signal 31 as a clock enable signal..

In one embodiment, the data value encoded in the first data frame 53a of the first data burst period 52a is always zero. This acts as a "start" bit to indicate the start of the data burst period 52a. When the control logic 30 detects the start data bit, for example by detecting the first falling edge of the signal occurring after time T_0 (i.e. at time T_1), this causes the control logic 30 to begin reading data. The data is read by detecting the state of the SysClkReq signal 31 during the data period 54 of each data frame 53 (following the first one, if the first data frame 53a is used to indicate the start of the data burst period 52a).

In embodiments in which the start bit 53a is used to signal the beginning of the data burst period 52a, asynchronous timing can be used for the data burst. In an alternative embodiment, the time period ΔT between the first rising edge of the SysClkReq signal 31 (at time T_0) and the start of the first data burst period 52a (at time T_1) is predetermined or selected in advance. In this way, the control logic 30 can begin reading data the time period ΔT after time T_0 . In this case, the first bit 53a can be used to represent data instead of indicating the start of the data burst period 52a.

10 The data period 54 and padding period 55 of each data frame 53a-z may last for any suitable time. For example, in one embodiment, the data period 54 and padding period 55 each last one clock cycle. In some embodiments, the duration of each padding period is greater than a certain fraction of the duration of each data period. Preferably, the duration of each padding period is at least as large as the duration of each data period. This ensures that the signal level averaged over a data burst period is greater than a threshold (even if all data bits are zero). For example, the threshold is the level such that an (analogue) signal having a level less than the threshold is considered to be in a low state while a signal having a level greater than the threshold is considered to be in a high state. In this way, an averaged signal can be used to represent a clock enable signal. Alternatively, or additionally, the timing of the data periods and padding periods may be selected so that changes in state of the SysClkReq signal 31 due to serial data is not interpreted by the first circuit 22 as changes in state of the clock enable component of the signal.

20 For example, each data period may be shorter than the period during which the SysClkReq signal 31 must remain in a particular state, after changing from one state to another, for the first circuit 22 to regard the clock enable component to have changed state. This ensures, for example, that a zero bit of data is not interpreted by the first circuit as the clock enable component changing to a low state. In order to enable the logic 34 to apply the correct timing for generation of the SysClkReq signal 31, the second ASIC 23 comprises an oscillator 36 which provides clock signals to the logic 34.

In one embodiment, the data burst period 52a comprises a pre-determined or otherwise selected number of data frames 53a-j, such as 8 or 16. This scheme limits the number of bits of data that can be communicated in a single data burst period 52a. However, in other embodiments, the data bit encoded
5 within the last data frame 53j of a data burst period 52a is used to indicate whether or not a further data burst 52b is to follow. For example, if the last data bit is equal to 1 (resulting from the signal being in a high state in the data period 54 of the last data frame 53j), then this signals that a further data burst period 52b will immediately follow the current one 52a. In this way a further
10 sequence of bits may be transmitted. If the data bit encoded in the last data frame 53z of a data burst period 52b is equal to 0 then this signals that no more data is to follow. This arrangement allows an arbitrary number of data burst periods to occur thereby making it possible to relay arbitrarily large amounts of data.

15
Following the final data burst period 52b, at a time T_2 , the SysClkReq signal 31 is maintained in a high state, shown as period 56 in Figure 4a. The high state of the SysClkReq signal 31 during period 56 indicates that at least one clock request signal 33a-c is still being asserted. The control logic 30 of the
20 first ASIC 22 therefore continues to provide clocking signals 24. In one embodiment, the data communicated to the first ASIC 22 in the SysClkReq signal 31 comprises information relating to which of the further ASICs 25a-c is requesting clocking signals 24a-c. Since the control logic 34 knows which of the clock request signals 33a-c are being asserted, the logic 34 of the second
25 ASIC 23 can encode this information into the SysClkReq signal 31 accordingly. Using this information, the control logic 30 of the first ASIC 22 may control the gates 28a-c such that, during period 56, respective clocking signals 24 are provided only to those further ASICs 25 that are requesting them. Clocking signals 24 are not provided to the other further ASICs 25. In
30 this way, the system 21 is dynamically configurable to enable optimisation of power consumption and performance. It is understood that other information may be communicated, such as data specifying which clocking mode to place the system 21 into.

At a time T_3 , none of the further ASICs 25a-c require clocking signals 24 and so no clock request signals 33 are asserted. During the subsequent period in which no clocking signals 24 are provided, shown as period 57 in Figure 4a, the system 21 is placed into sleep mode once again. At time T_3 , the logic 34
5 of the second ASIC 23 detects that no clock request signals 33 are being asserted and places the SysClkReq signal 31 in a low state. When the control logic 30 receives the SysClkReq signal 31 in a low state this causes the control logic 30 to initiate the necessary procedure to enter sleep mode.

10 In one embodiment, the control logic 30 of the first ASIC 22 is arranged so that the SysClkReq signal 31 must remain in a low state for a certain time period before initiating entry into sleep mode. This ensures, for example, that sleep mode is not entered erroneously due to fluctuations in the SysClkReq signal 31 during the data burst periods 52a, 52b. In one embodiment, the
15 SysClkReq signal 31 must be maintained in a low state for a period at least equal to the duration of one or more data frames 53, or at least greater than the duration of one data period, for the control logic 30 to initiate entry into sleep mode. If this condition is imposed, since the SysClkReq signal 31 goes high during each data frame 53 (at least during the padding period 55), and is
20 not maintained in a low state for longer than one data period, it is not possible for the system 21 to enter sleep mode during a data burst period 52a, 52b. Low signal pulses during a data burst period would not then be interpreted by the first circuit 22 as the clock enable component changing to a low state.

25 In order that the SysClkReq signal can function as both a clock enable signal and data carrier, the control logic 30 comprises means to enable the different components (clock enable and data components) of the SysClkReq signal 31 to be separated or distinguished. For example, in the illustrated embodiment, to separate the clock enable component and data component of the
30 SysClkReq signal, a low-pass filtered 39 and non-filtered 40 version of the SysClkReq signal are used. To produce the low-pass filtered version 39, the SysClkReq signal 31 is first passed through a digital to analogue converter 37 and then through a low-pass filter 38. This process removes the high-rate fluctuations of the SysClkReq signal 31 during the data burst periods 52a,

52b. The data component of the SysClkReq signal 31 is thus filtered out, leaving the clock enable component. The resulting low-pass filtered signal 39 has an overall form that is similar to the clock enable signal 11 illustrated in Figure 2, except that the rising and falling edges rise and fall more gradual in the filtered SysClkReq signal 39. The filtered SysClkReq signal 39 can then be used substantially in the same way as the clock enable signal 11 shown in Figure 2.

The delay, ΔT between the first rising edge of the SysClkReq signal 31 at time T_0 and the beginning of the first data burst period 52a at time T_1 provides time for the low pass filtered signal 39 to rise from the low state to the high state. Also, in this example the SysClkReq signal 31 is in a high state for at least half of every data burst period 52 (due to the inclusion of the padding periods 55). This means that the low-pass filtered signal 39 will remain substantially at the same level throughout the data burst period 52a, 52b, or will at least remain above the threshold level required for the signal to be regarded as being in a high state. If the padding periods 55 were not included then a sequence of substantially zero data bits may cause the low pass filtered signal to gradually decrease to a level corresponding to the low state if the data burst period 52a, 52b is sufficiently long. If the length of a continuous data burst period 52a, 52b is sufficiently short, then the padding periods may not be required since the data burst may terminate before the fluctuations in the signal during the data burst period had a chance to influence the filtered signal.

While a low-pass filtered version of the SysClkReq signal 39 is used as a clock enable signal, a non-filtered version of the SysClkReq signal 40 is used to extract the data encoded in the data burst periods 52a, 52b by reading the state of the signal during each data period 54 in the data burst periods 52a, 52b.

The clock enable component and the data component in the SysClkReq signal 31 may be separated or distinguished in other ways. For example, the

distinguishing of the clock enable component and data component in the SysClkReq signal 31 may be performed entirely digitally, without requiring a low-pass filter or other analogue components. In one embodiment, when the SysClkReq signal 31 is used as a clock enable signal, digital processing
5 circuitry comprised in the control logic 30 of the first ASIC 22 may be arranged to ignore the state of the SysClkReq signal 31 during the data periods 54. This is possible if the control logic is able to determine the beginning of a data burst period 52a (indicated, for example, by the first bit of the data burst) and has information relating to the timing of the data periods and padding periods
10 (which may be predetermined or otherwise specified in advance). Conversely, when the SysClkReq signal 31 is used as a data transmitting signal, the control logic uses the state of the signal 31 within the data periods 54 to read the data. In one embodiment the SysClkReq signal 31 is successively delayed using a chain of D-flip-flop components. The resulting delayed data bits are
15 then combined using OR logic gates. If one or more of the delayed data bits has a value of 1 then the first circuit 22 interprets this as the clock enable component being in a high state. As with analogue methods, when digital processing is used to distinguish the clock enable and data components of the SysClkReq signal 31, the presence of the padding periods 55 and/or the
20 delay, ΔT 51, mentioned above, facilitates the process of separating or distinguishing these components.

Since the SysClkReq signal 31 is transmitted from the second ASIC 23 to the first ASIC 22 on a single line there is no need for additional pins on the first
25 ASIC 22 to accommodate multiple transmission lines. Instead, a single signal transmitted on a single line provides dual function of a clock enable signal together with data transfer. Since the SysClkReq signal 31 functions as a clock enable signal, the embodiments described above are compatible with existing signalling techniques and protocols. Furthermore, the ability to also
30 transmit data using the SysClkReq signal 31 enhances functionality. These advantages are realised with only small modifications of existing designs.

- Thus, in some embodiments, it is advantageous for an RF ASIC to have information regarding which components of the system requested clocking signals, causing the sleep mode to be exited. This allows optimisation of power management by providing clocking signals only to those components
- 5 which actually requested them. This may be achieved without requiring additional pins, or limited additional pins. Additional information may be communicated while supporting existing signalling techniques for back-compatibility.
- 10 The techniques of the present invention are not limited to the specific embodiments described above. For example, information other than, or in addition to, clocking information may be transmitted. A signal other than a clock enable signal may be used to transmit the serial data. The techniques described above may be used to transmit data between ASICs other than PM
- 15 ASICs and RF ASICs. More generally, a signal used in a system having at least two states may be modulated to encode data into the signal. The two signal states may be used for example to cause enabling or disabling of one or more functions in the system, or to cause two respective functions to be performed. The signal is modulated to encode the data while the signal is
- 20 being maintained in one of the states. The frequency at which the signal changes from one state to the other (for the purpose of causing enabling/disabling or performance of functions) will therefore be lower than the frequency of the encoded data. The modulation is performed such that the encoded data is distinguishable from the signal. Other modifications will
- 25 readily occur to those skilled in the art.

Claims

1. A first circuit (23) for transmitting information to a second circuit (22),
the first circuit (23) comprising a signal generator (34) arranged to
5 generate a signal (31) for transmission to the second circuit (22), the
signal (31) having a first state for causing the second circuit (22) to
perform a first operation and having a second state for causing the
second circuit (22) to perform a second operation; wherein the signal
generator (34) comprises means to modulate the generated signal (31)
10 during a data burst period (52a, 52b) to encode data into the generated
signal (31), the data burst period (52a, 52b) occurring within a period in
which the generated signal (31) is in the second state, the signal (31)
being modulated such that the encoded data is distinguishable from the
generated signal (31).
15
2. A circuit according to claim 1 in which the data burst period (52a, 52b)
comprises a plurality of data periods (54) during which the state of the
signal is representative of the data encoded in the signal (31).
- 20 3. A circuit according to claim 2 in which the data periods (54) are
separated by padding periods (55) during which the signal (31) is in the
second state.
4. A circuit according to claim 3 in which each padding period (55) is at
25 least as large as the duration of each data period (54).
5. A circuit according to any of claims 2 to 4 in which the signal (31) is in
the first state during the first data period.
- 30 6. A circuit according to any of claims 2 to 5 in which the data burst period
(52a, 52b) comprises a predetermined number of data periods (54).

7. A circuit according to claim 6 in which the state of the signal (31) during the last data period is indicative of whether a further data burst (52b) period will follow the current one (52a).
- 5 8. A circuit according to any preceding claim in which the data burst period is initiated following a delay after the signal changes from being in the first state to being in the second state.
9. A circuit according to claim 1 in which the signal (31) is modulated such
10 that the time-averaged level of the modulated signal when the generated signal (31) is in the second state is greater than a threshold.
10. A circuit according to any preceding claim in which the first operation is
15 restricting the generation of clocking signals (24a-c) by the second circuit (22) and the second operation is providing one or more clocking signals (24a-c) by the second circuit. (22)
11. A circuit according to claim 10 in which the data encoded into the signal
20 (31) includes data identifying which components of the system (21) are requesting clocking signals (24a-c).
12. A circuit according to any preceding claim in which the system (21) is a wireless cellular system.
- 25 13. A system (21) for signalling information between a first circuit (23) and a second circuit (22), the first circuit (23) comprising a circuit according any preceding claim, the second circuit (22) comprising an input to receive the signal (31) generated by the first circuit (23).
- 30 14. A system according to claim 13 in which the second circuit (22) comprises means for separating the encoded data from the signal (31).

15. A system according to claim 14 in which the means for separating the encoded data from the signal (31) comprises one or more of: a low pass filter; and a digital processing circuit.

1/3




Figure 1

2/3




Figure 3




Figure 2




Figure 4a




Figure 4b

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2009/067814

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L7/02 G06F1/32
ADD. H04L7/033

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 860 815 A1 (HIGH TECH COMP CORP [TW]) 28 November 2007 (2007-11-28) abstract; claims 1,6; figures 1-3 paragraphs [0002], [0005] - [0007], [0015] - [0017] paragraphs [0025], [0026], [0028]	1-15
X	EP 0 280 391 A2 (HEWLETT PACKARD CO [US]) 31 August 1988 (1988-08-31) abstract; claims 1,2,5,6; figures 1a,1b column 1, line 51 - column 2, line 13 column 3, line 11 - column 4, line 15 ----- -/--	1-15

☒ Further documents are listed in the continuation of Box C.

☒ See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier document but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
"&" document member of the same patent family

Date of the actual completion of the international search

11 March 2010

Date of mailing of the international search report

25/03/2010

Name and mailing address of the ISA/
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Wolters, Robert

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2009/067814

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 864 960 A1 (NOKIA MOBILE PHONES LTD [FI] NOKIA CORP [FI]) 16 September 1998 (1998-09-16) abstract; claims 1,2,5-13; figures 3,4 column 1, line 21 - column 2, line 25 column 3, line 55 - column 4, line 2 column 4, line 18 - column 5, line 8 -----	1-15
X	US 2006/047988 A1 (OKADA TAKEYUKI [JP]) 2 March 2006 (2006-03-02) abstract; claims 1-4; figures 1,2 paragraphs [0010] - [0012] paragraphs [0019] - [0025], [0030] - [0035], [0038] - [0040] -----	1-15
A	EP 1 819 048 A1 (FUJITSU LTD [JP]) 15 August 2007 (2007-08-15) abstract; claims 1,2; figures 3,4 paragraphs [0011] - [0019] -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2009/067814

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1860815	A1	28-11-2007	AT 429096 T 15-05-2009 CN 101079017 A 28-11-2007 ES 2324121 T3 30-07-2009
EP 0280391	A2	31-08-1988	DE 3889944 D1 14-07-1994 DE 3889944 T2 15-09-1994 JP 2708170 B2 04-02-1998 JP 63226154 A 20-09-1988 US 4782481 A 01-11-1988
EP 0864960	A1	16-09-1998	DE 69825344 D1 09-09-2004 DE 69825344 T2 21-07-2005 GB 2323188 A 16-09-1998 JP 11024928 A 29-01-1999 JP 2005196811 A 21-07-2005 US 6202163 B1 13-03-2001
US 2006047988	A1	02-03-2006	JP 2006065471 A 09-03-2006
EP 1819048	A1	15-08-2007	CN 101073204 A 14-11-2007 WO 2006059379 A1 08-06-2006 JP 4382096 B2 09-12-2009 US 2007242556 A1 18-10-2007