US 20060114836A1

(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2006/0114836 A1

Pollin et al. (43) Pub. Date: Jun. 1, 2006

(54) METHOD FOR OPERATING A COMBINED MULTIMEDIA -TELECOM SYSTEM

(76) Inventors: **Sofie Pollin**, (US); **Bruno Bougard**, (US); **Gregory Lenoir**, (US); **Francky Catthoor**, (US)

Correspondence Address:
KNOBBE MARTENS OLSON & BEAR LLP
2040 MAIN STREET
FOURTEENTH FLOOR
IRVINE, CA 92614 (US)

(21) Appl. No.: **11/247,403**

(22) Filed: **Oct. 11, 2005**

Related U.S. Application Data

(63) Continuation of application No. 10/922,371, filed on Aug. 20, 2004.

(60) Provisional application No. 60/617,897, filed on Oct. 12, 2004.

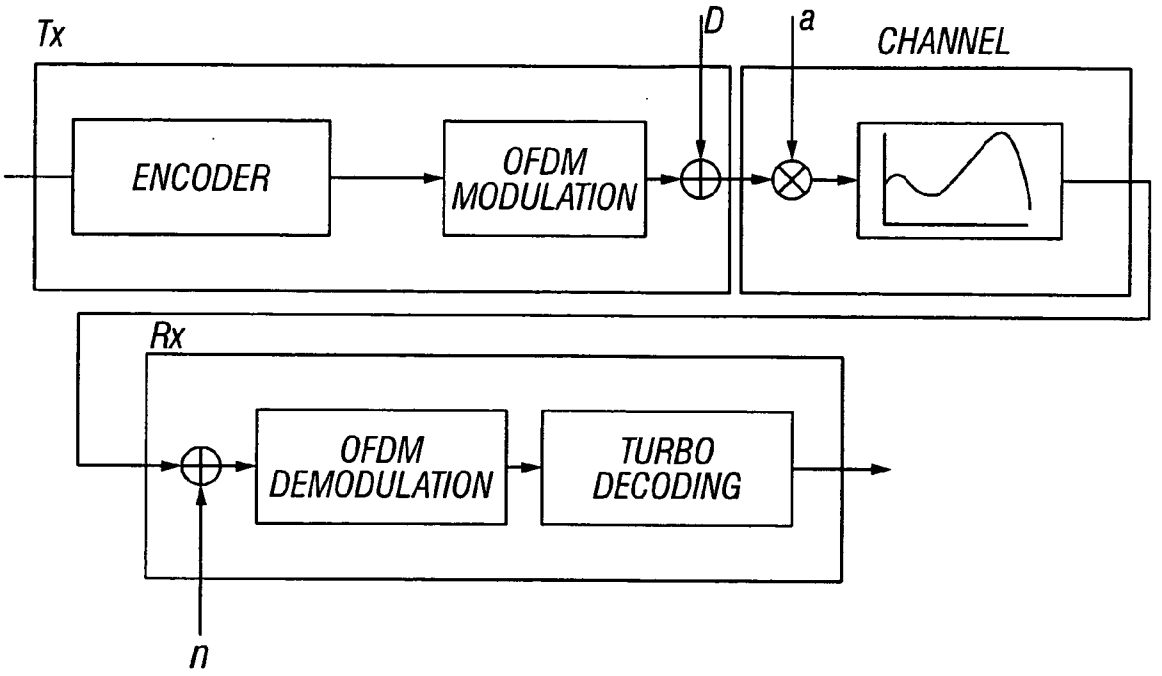
Publication Classification

(51) **Int. Cl.**
H04L 12/26 (2006.01)

(52) **U.S. Cl.** **370/252**

(57) **ABSTRACT**

Presented is a method of managing the operation of a system including a processing subsystem configured to run a multimedia application and a telecommunication subsystem. The method includes determining telecom environment conditions, and selecting a working point from a plurality of predetermined working points. The selecting is based at least in part on the determined environmental conditions. The method also includes setting control parameters in the multimedia application and/or the telecommunication subsystem to configure the system to operate at the selected working point, and operating the system at the selected working point.



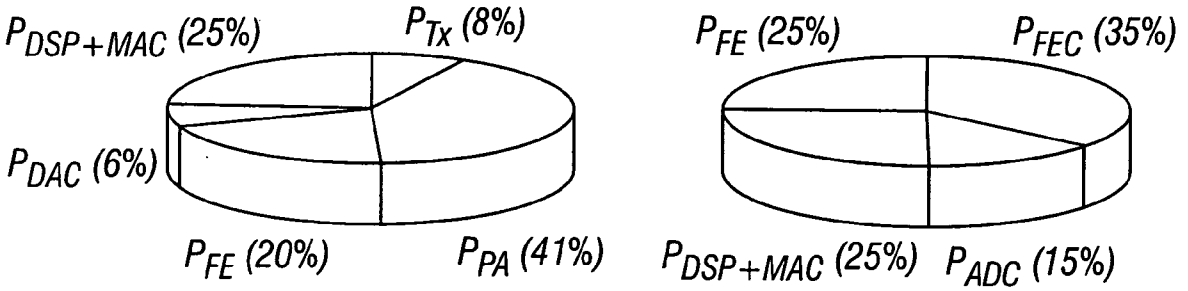


FIG. 1

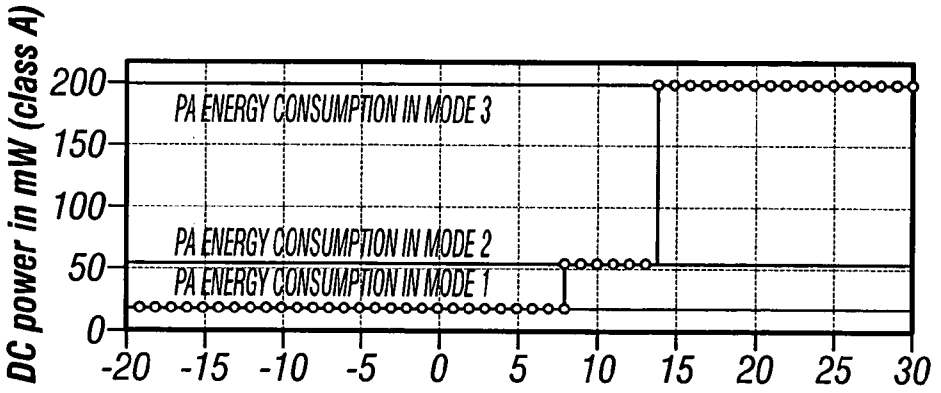


FIG. 2A

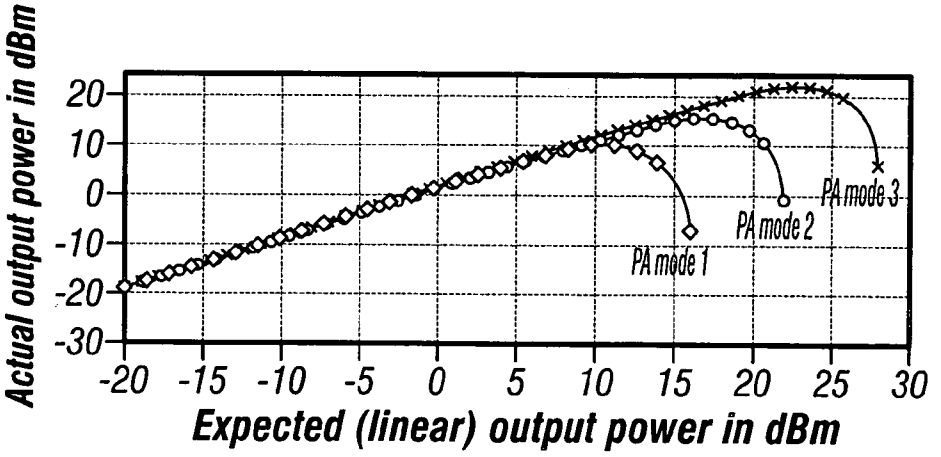


FIG. 2B

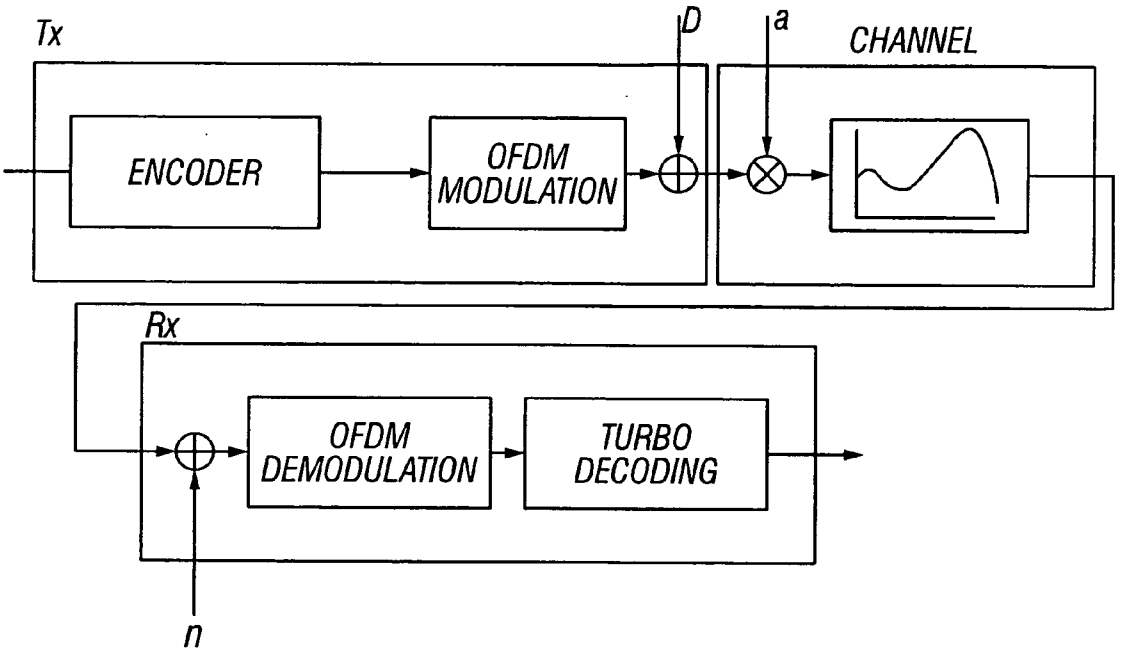


FIG. 3

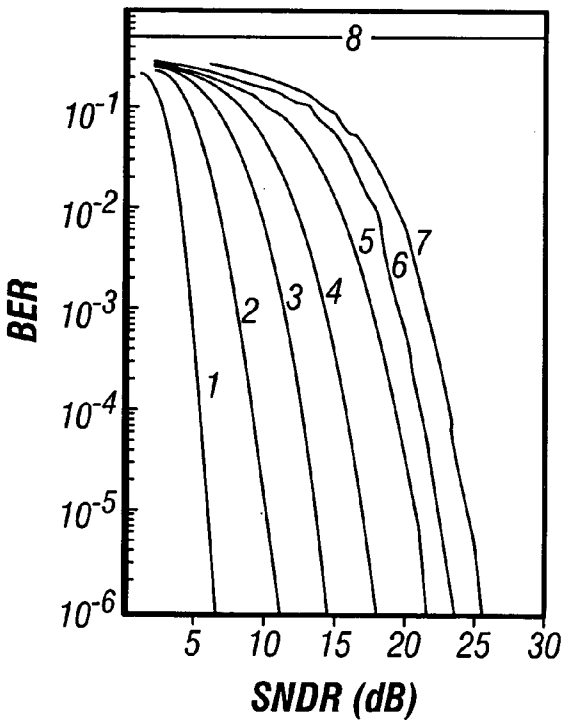


FIG. 4A

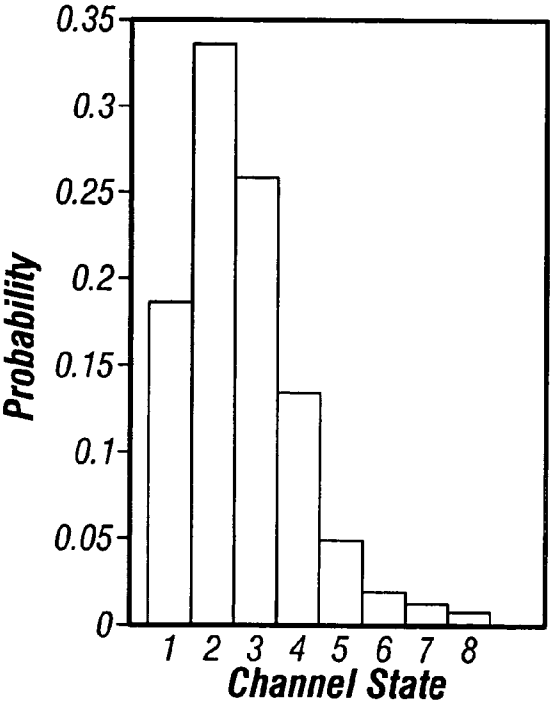


FIG. 4B

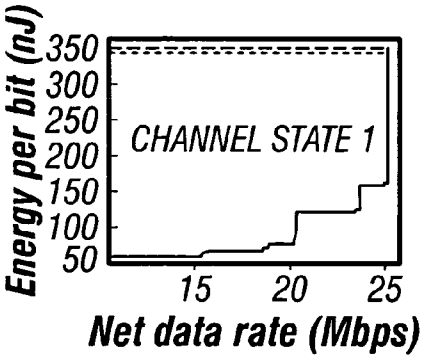


FIG. 5A

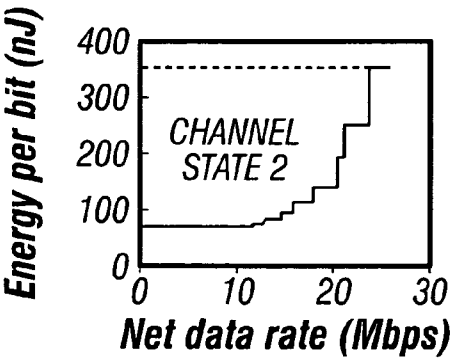


FIG. 5B

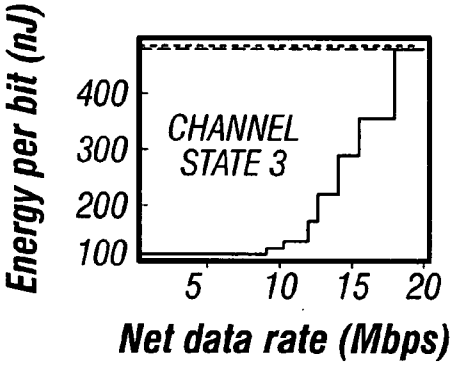


FIG. 5C

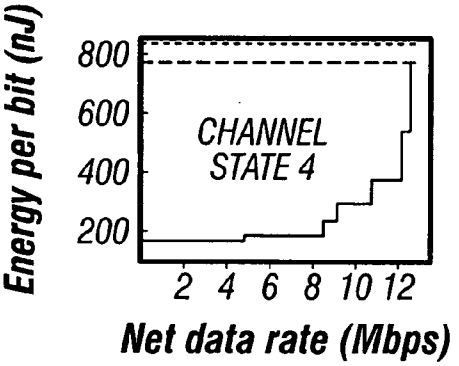


FIG. 5D

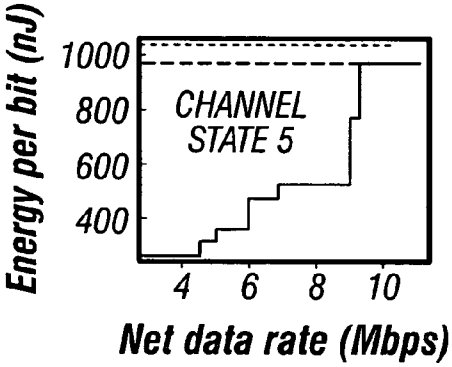


FIG. 5E

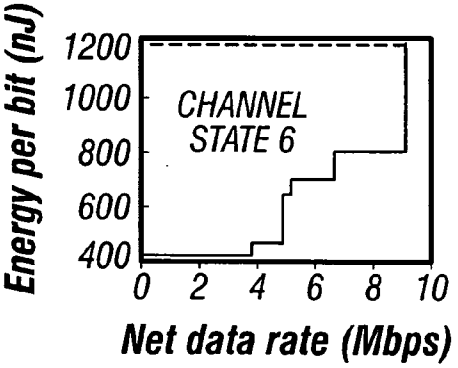
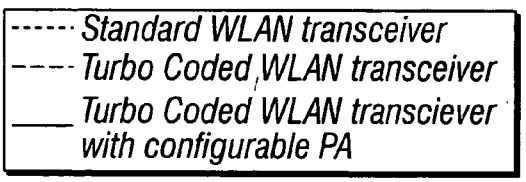


FIG. 5F



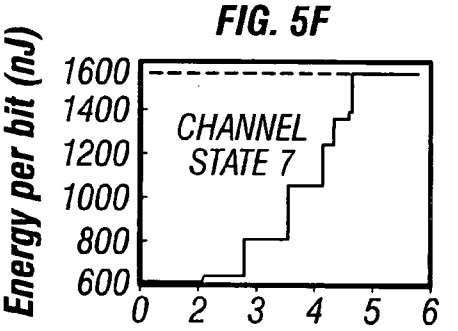


FIG. 5G

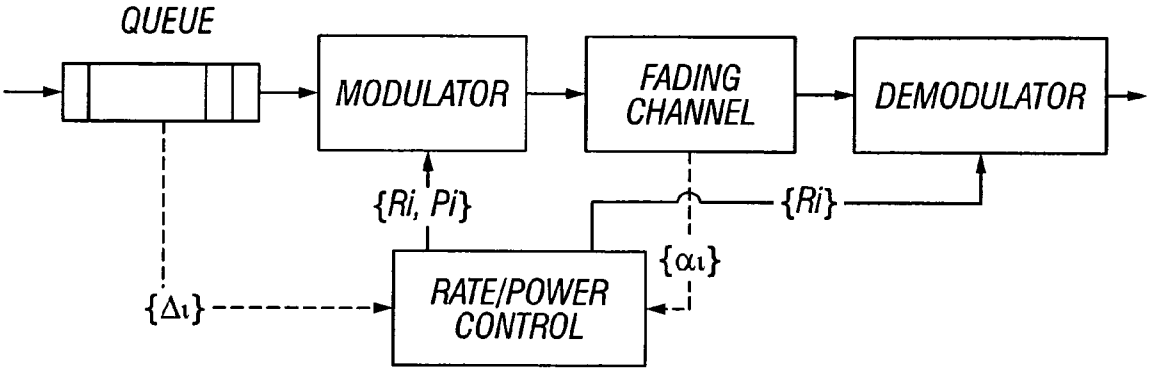

QUEUE
MODULATOR
FADING CHANNEL
DEMODULATOR
{Ri, Pi}
{Ri}
{αι}
{Δι}
RATE/POWER CONTROL

FIG. 6

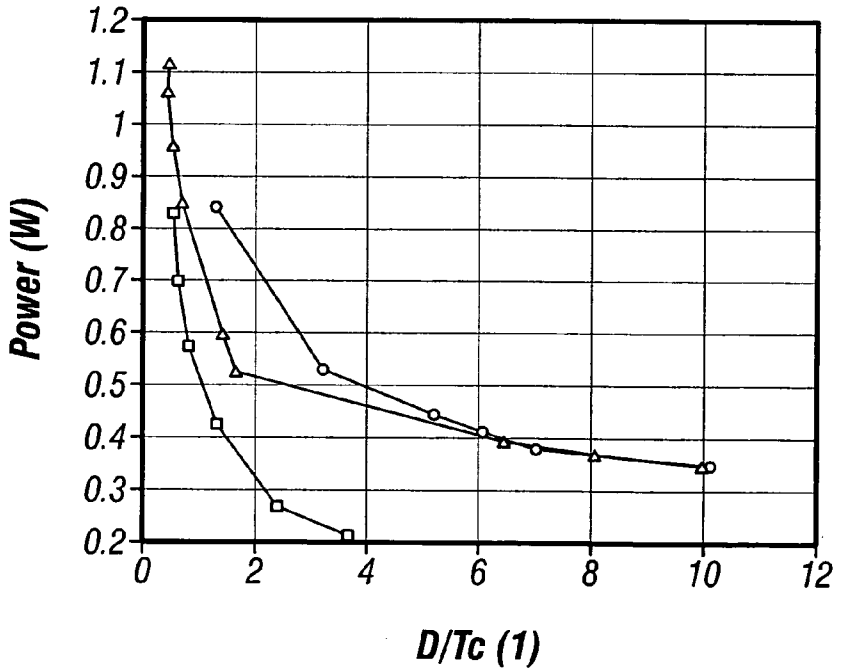

Power (W)
1.2
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0
2
4
6
8
10
12
D/Tc (1)

FIG. 7

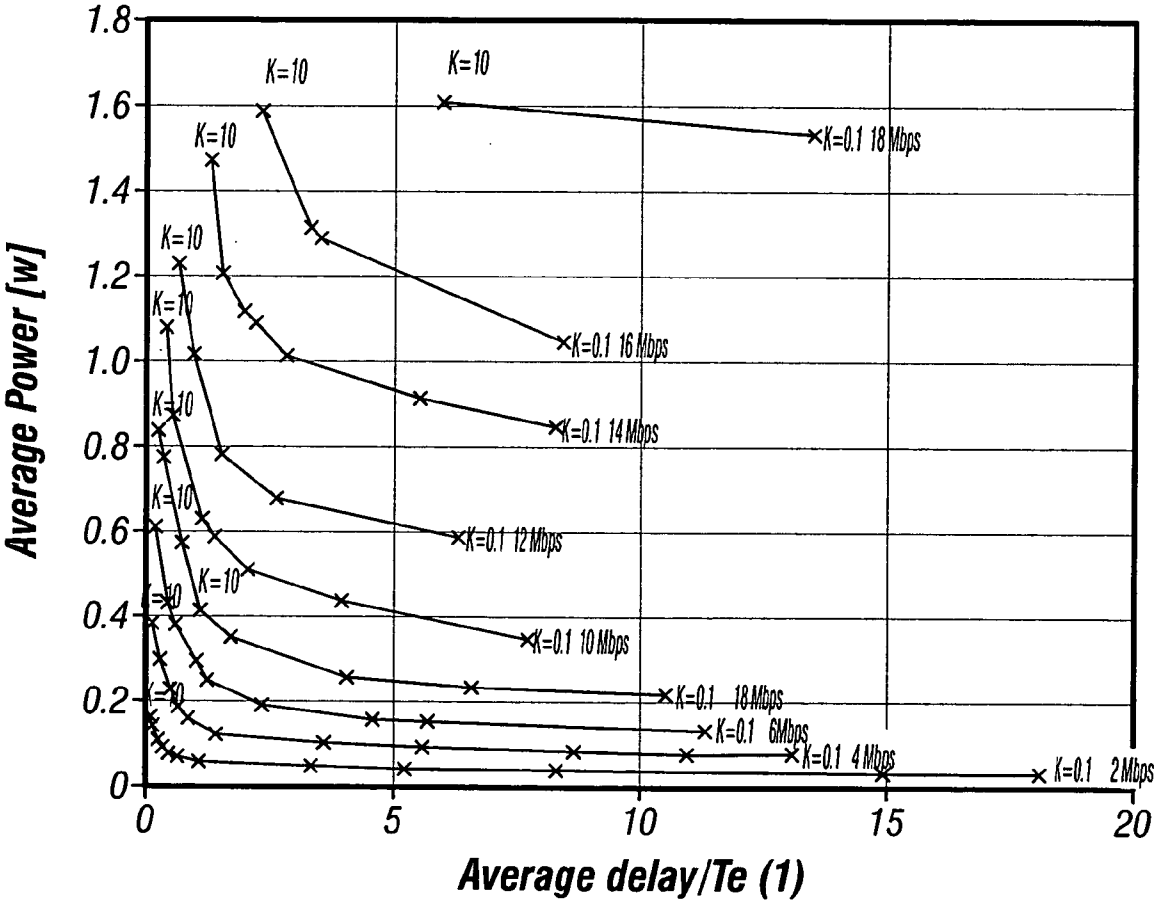

Average Power [w]
1.8
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0
0
5
10
15
20
Average delay/Te (1)
K=10
K=0.1 18 Mbps
K=0.1 16 Mbps
K=0.1 14 Mbps
K=0.1 12 Mbps
K=0.1 10 Mbps
K=0.1 8 Mbps
K=0.1 6Mbps
K=0.1 4 Mbps
K=0.1 2 Mbps

FIG. 8

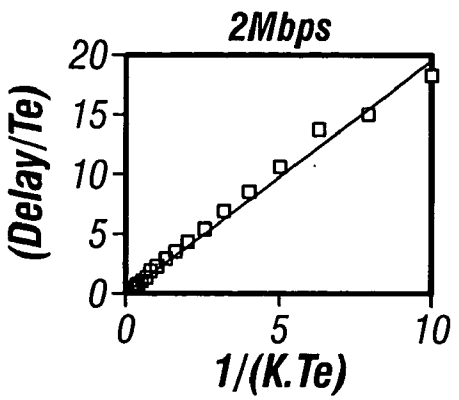


FIG. 9A

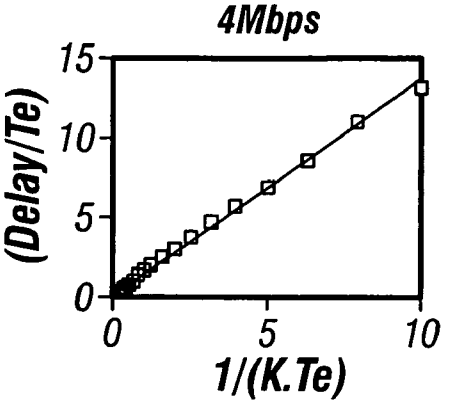


FIG. 9B

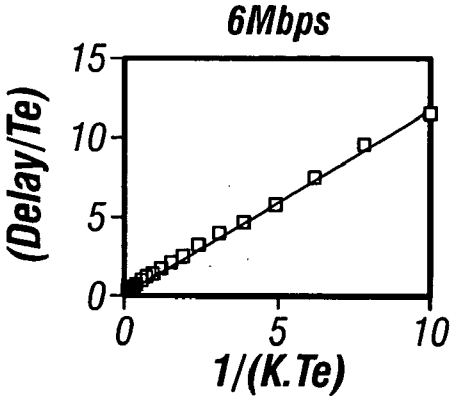


FIG. 9C

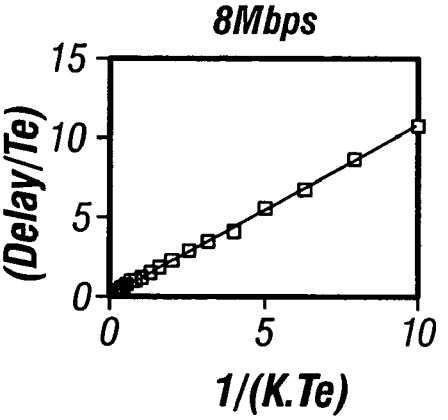


FIG. 9D

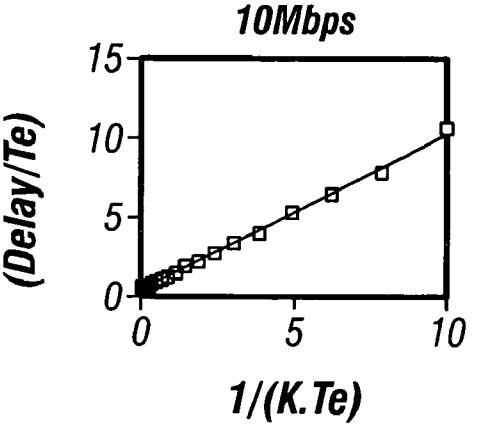


FIG. 9E

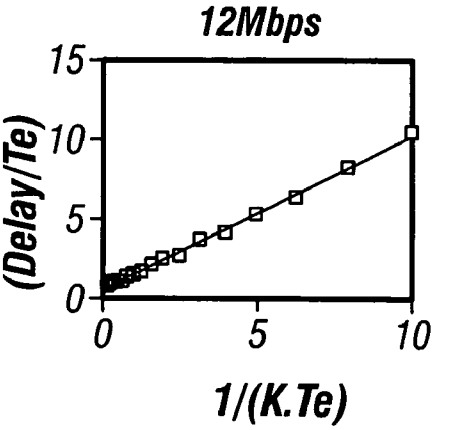


FIG. 9F

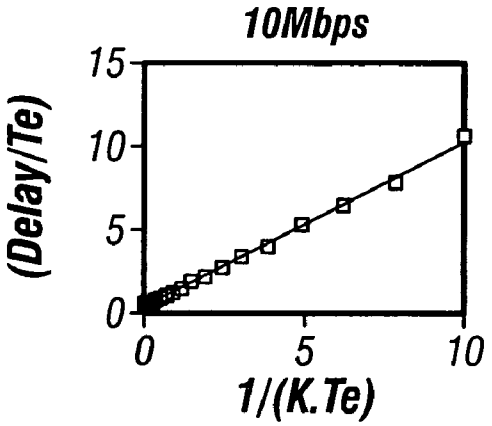


FIG. 9G

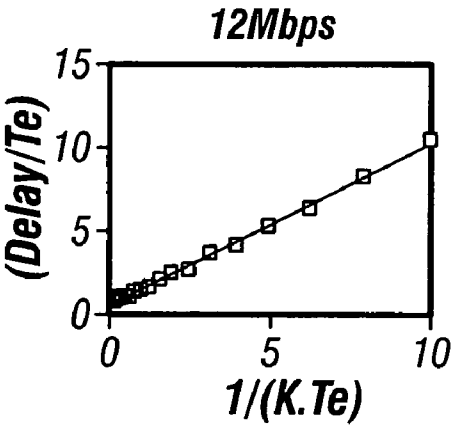


FIG. 9H

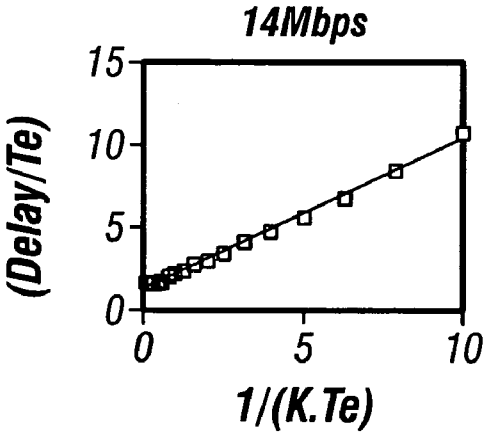


FIG. 9I

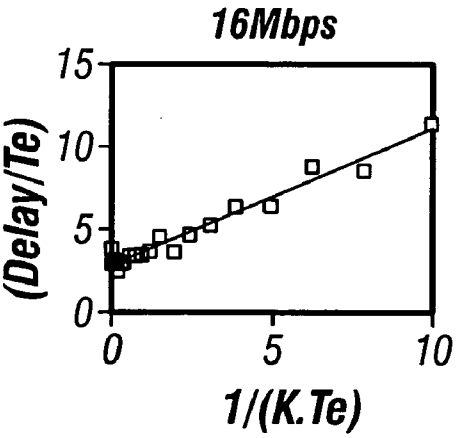


FIG. 9J

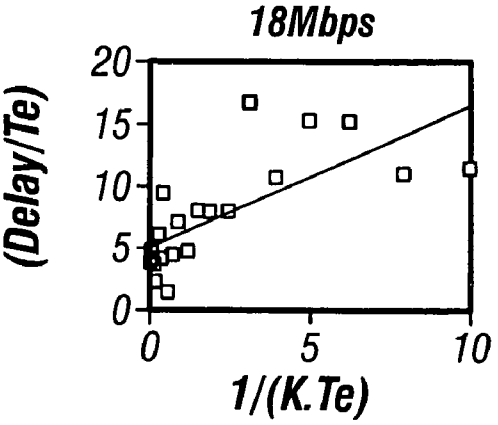


FIG. 9K

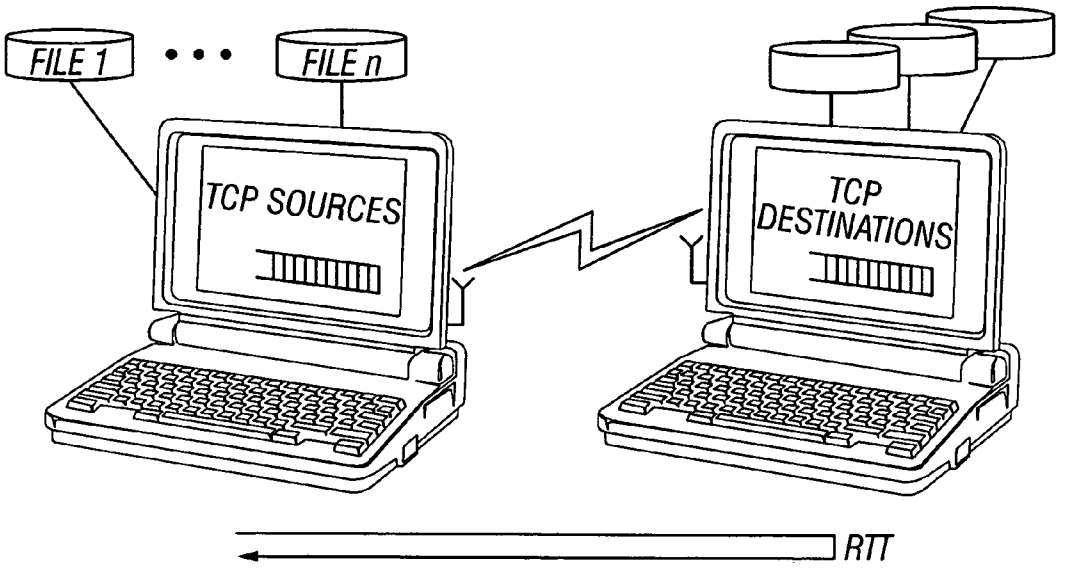

FILE 1
FILE n
TCP SOURCES
TCP DESTINATIONS
RTT

FIG. 10

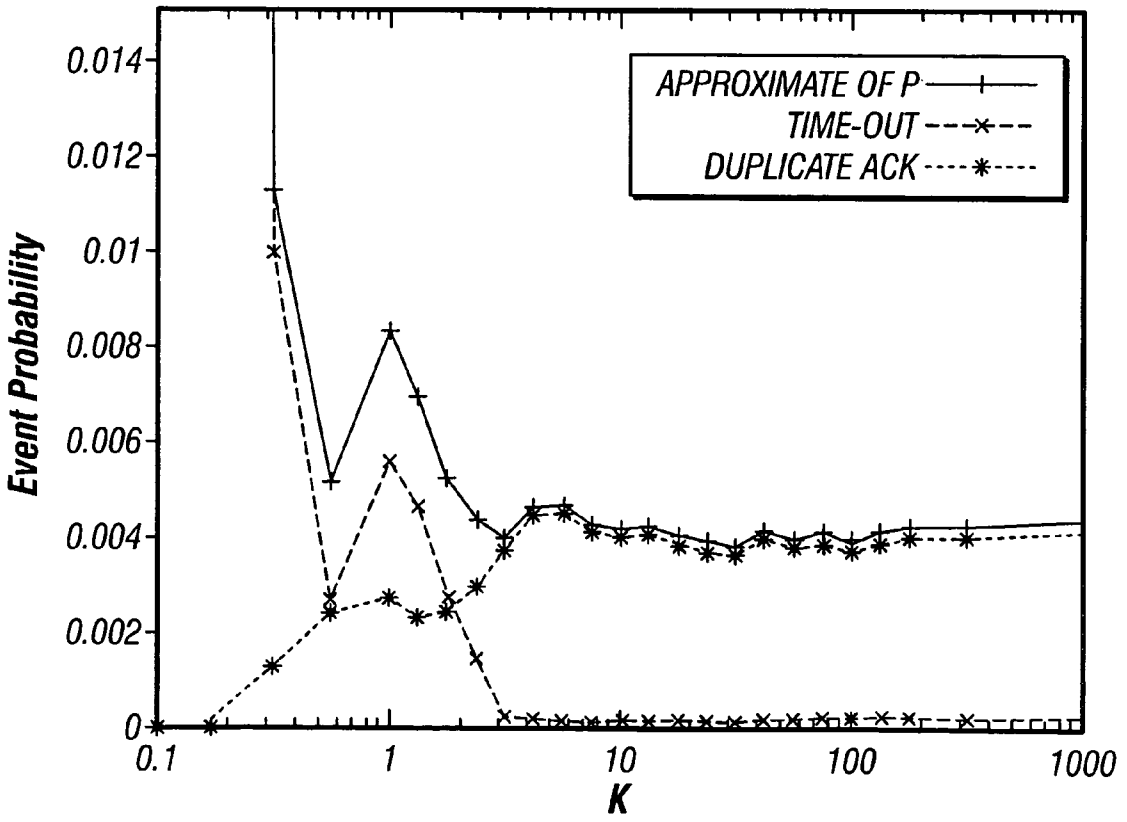

APPROXIMATE OF P
TIME-OUT
DUPLICATE ACK
Event Probability
0.014
0.012
0.01
0.008
0.006
0.004
0.002
0
0.1
1
10
100
1000
K

FIG. 11

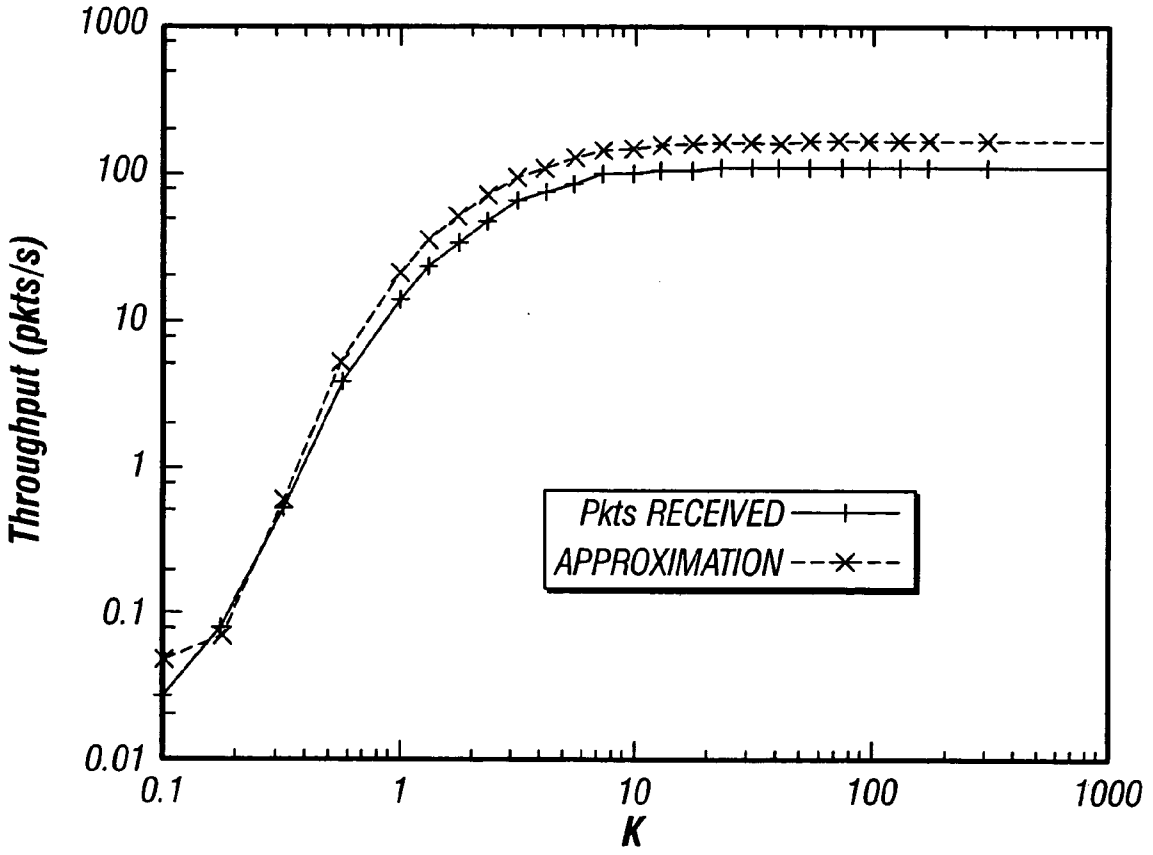

1000
100
10
1
0.1
0.01
Throughput (pkts/s)
0.1
1
10
100
1000
K
Pkts RECEIVED
APPROXIMATION

FIG. 12

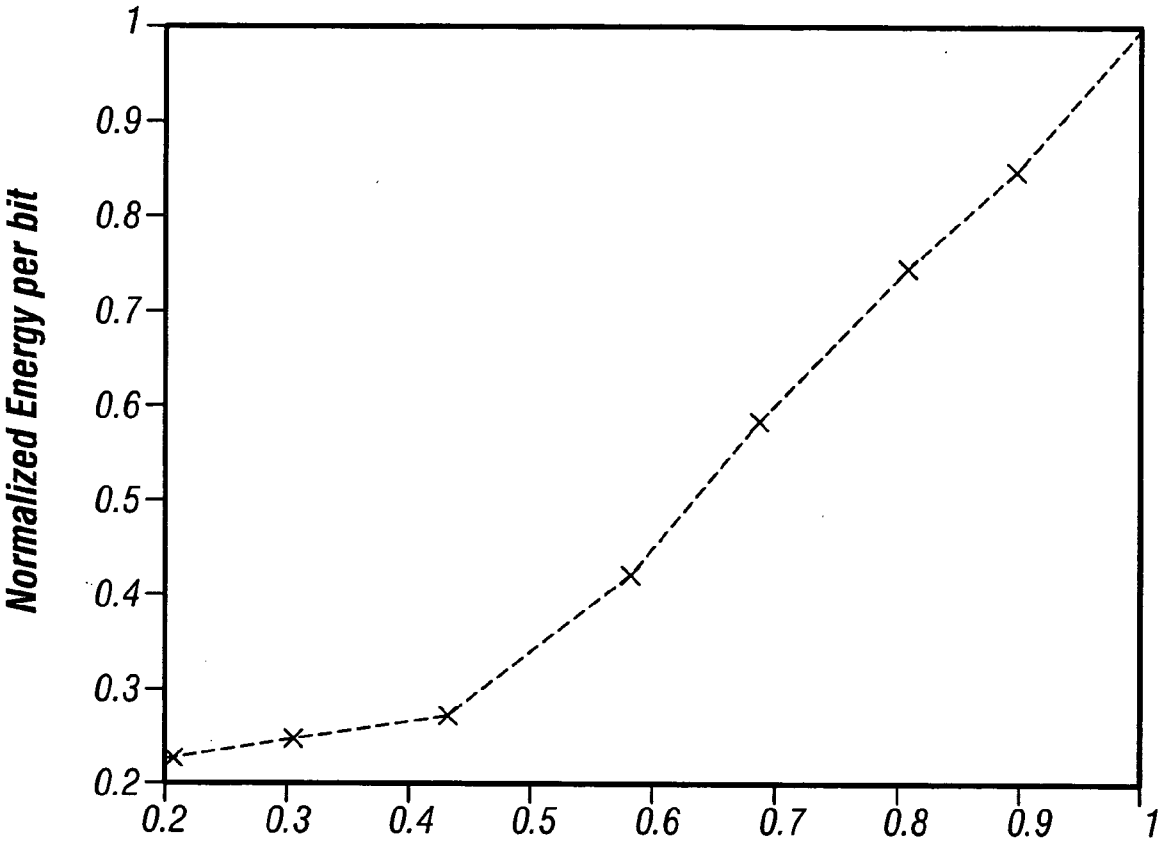

Normalized Energy per bit
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Normalized Throughput

FIG. 13

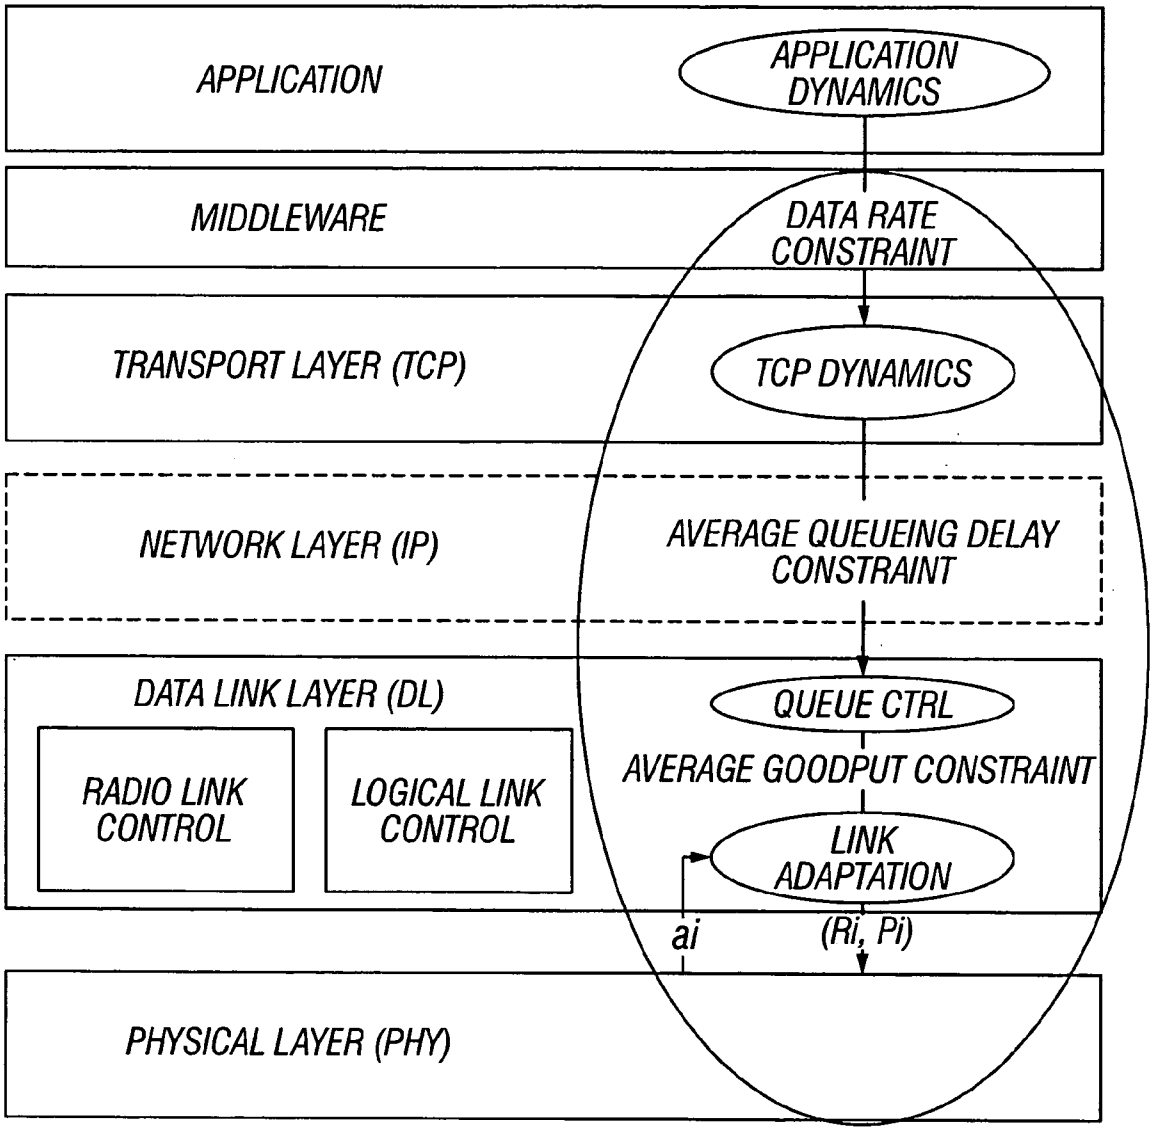

APPLICATION
APPLICATION DYNAMICS
MIDDLEWARE
DATA RATE CONSTRAINT
TRANSPORT LAYER (TCP)
TCP DYNAMICS
NETWORK LAYER (IP)
AVERAGE QUEUEING DELAY CONSTRAINT
DATA LINK LAYER (DL)
QUEUE CTRL
RADIO LINK CONTROL
LOGICAL LINK CONTROL
AVERAGE GOODPUT CONSTRAINT
LINK ADAPTATION
ai
(Ri, Pi)
PHYSICAL LAYER (PHY)

FIG. 14

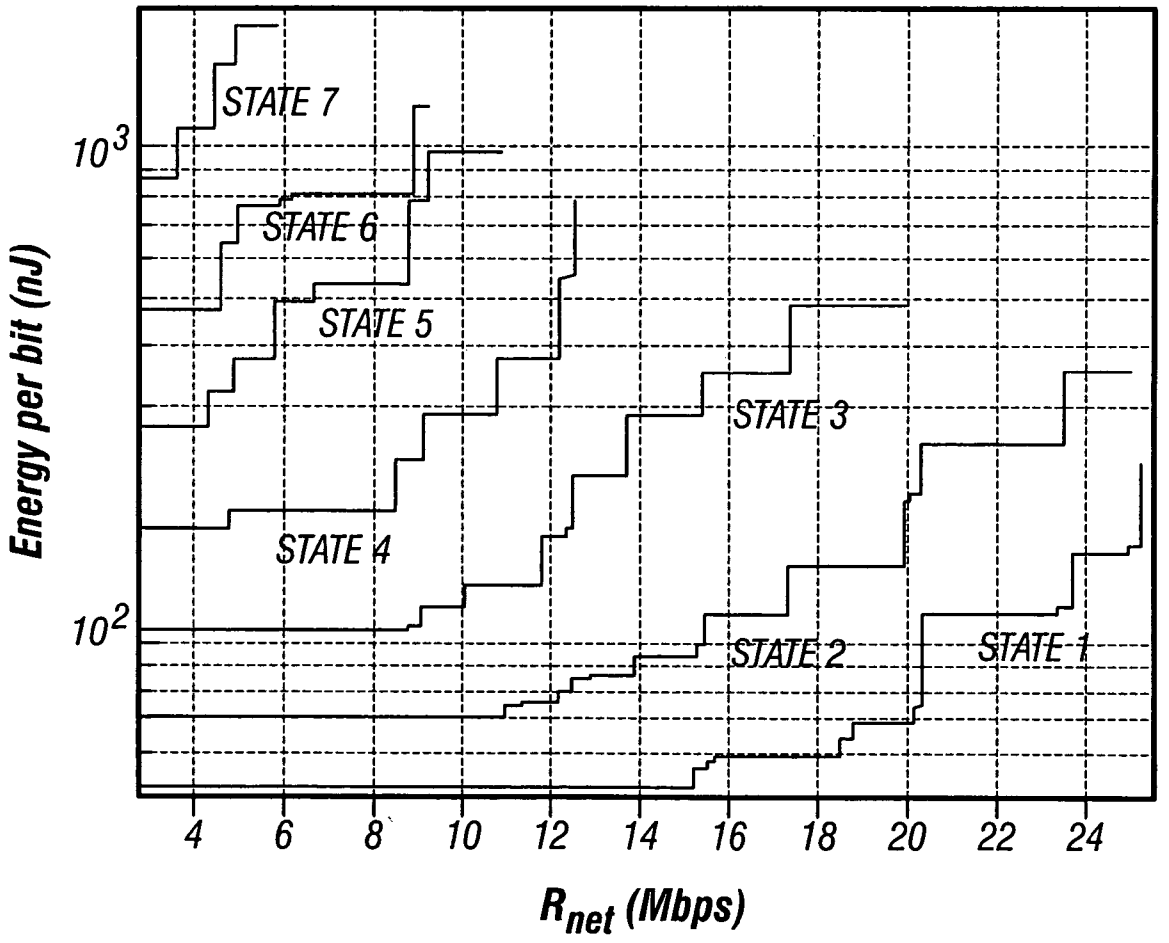

STATE 7
STATE 6
STATE 5
STATE 4
STATE 3
STATE 2
STATE 1
10^3
10^2
Energy per bit (nJ)
4
6
8
10
12
14
16
18
20
22
24
R_{net} (Mbps)

FIG. 15

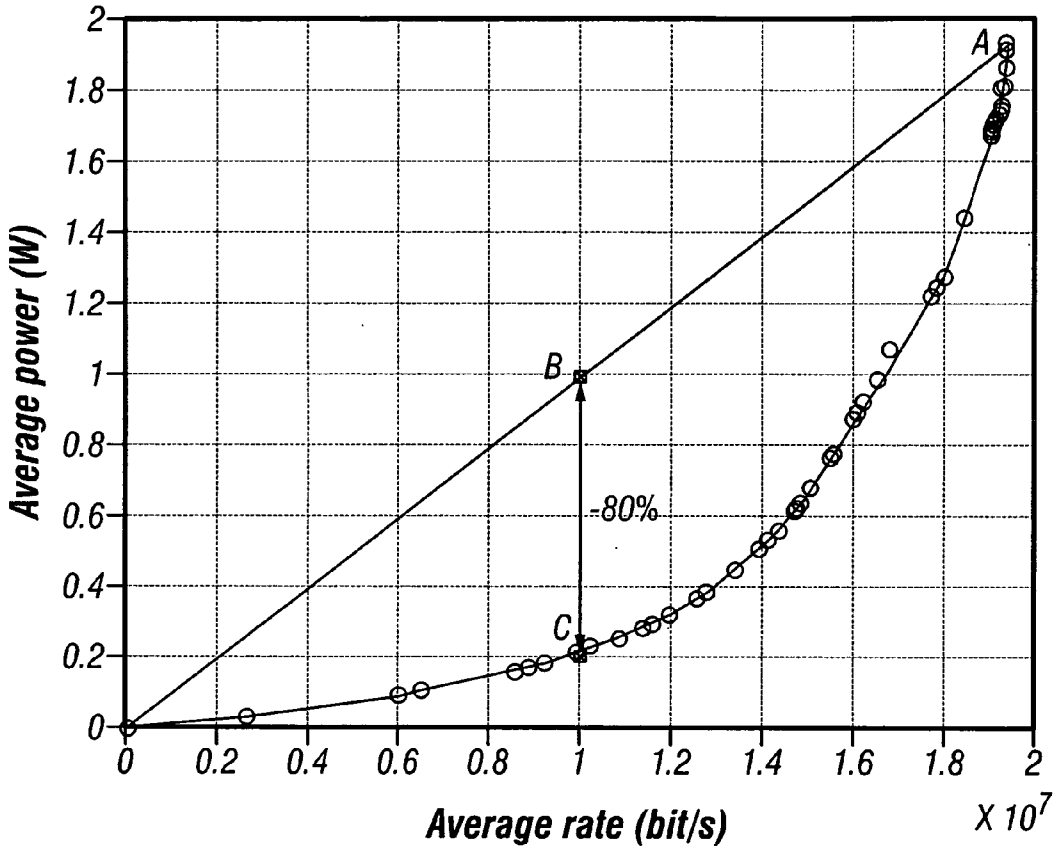

Average power (W)
Average rate (bit/s)
X 10^7
A
B
C
-80%
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2

FIG. 16

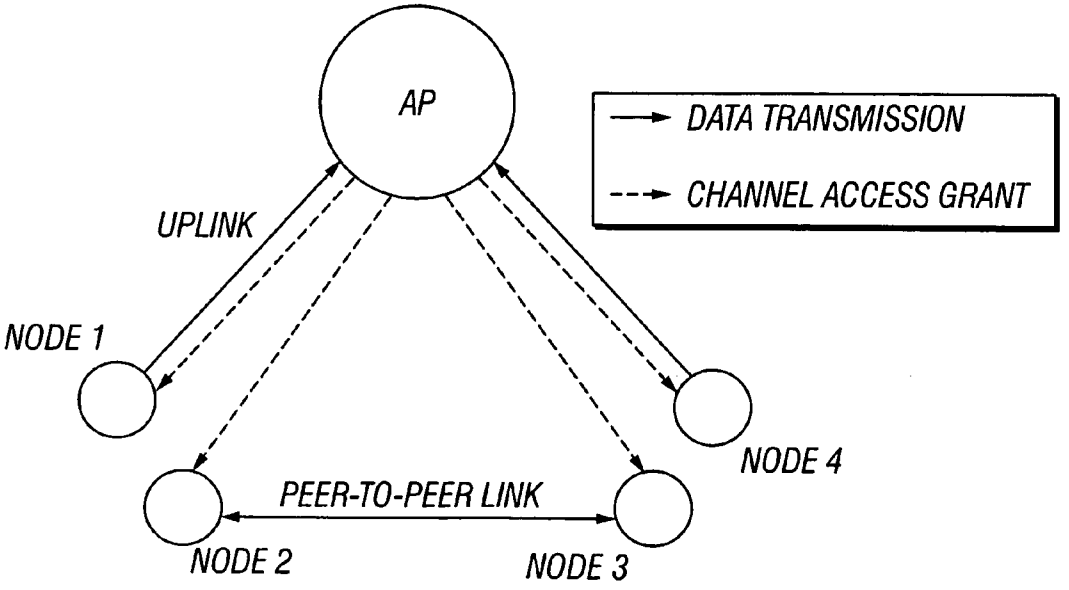

AP
DATA TRANSMISSION
CHANNEL ACCESS GRANT
UPLINK
NODE 1
NODE 2
NODE 3
NODE 4
PEER-TO-PEER LINK

FIG. 17

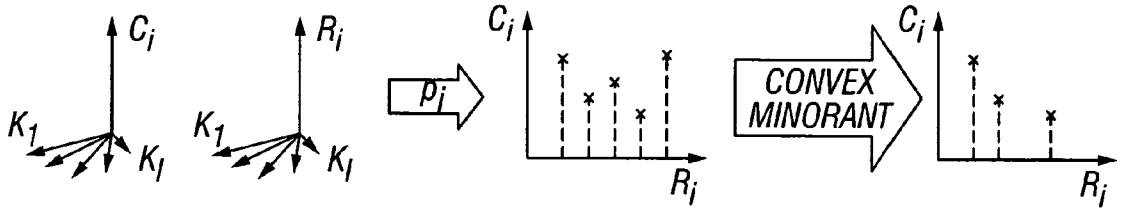


FIG. 18

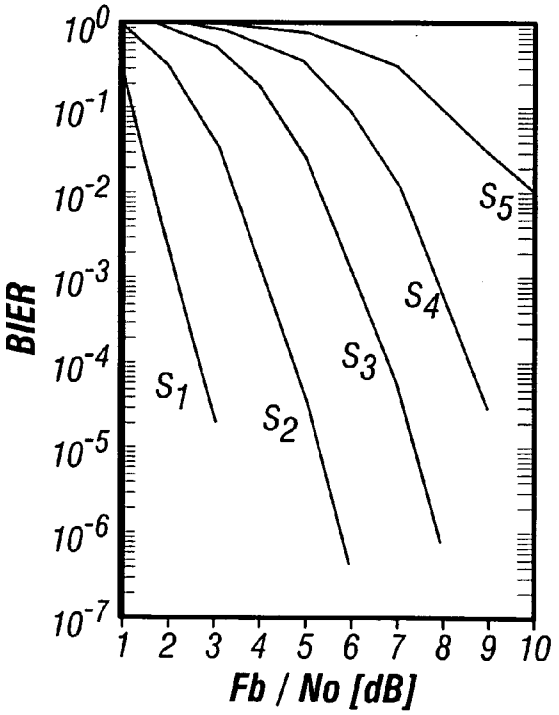


FIG. 19A

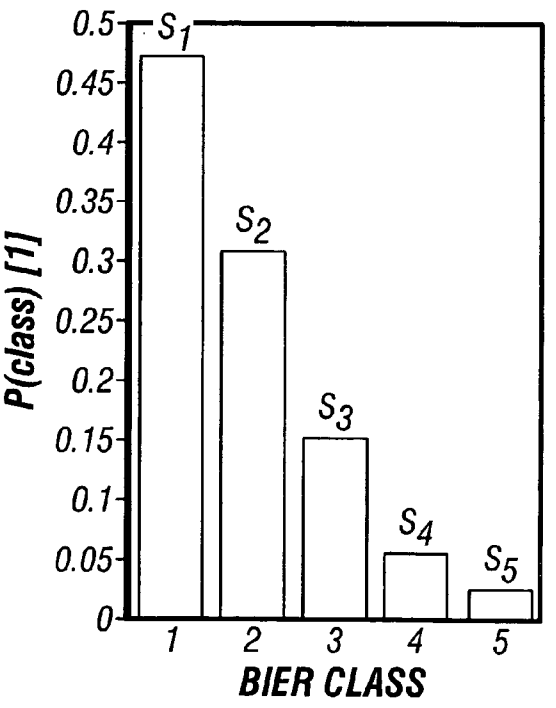


FIG. 19B

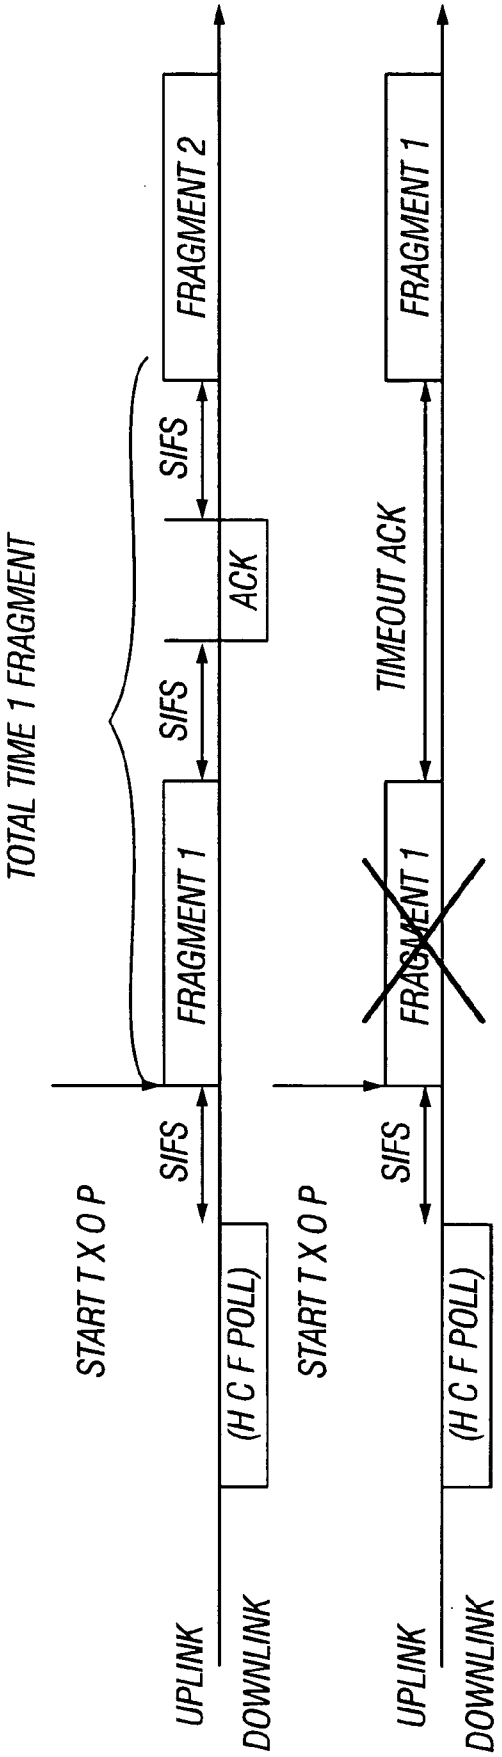

TOTAL TIME 1 FRAGMENT
START T X O P
SIFS
FRAGMENT 1
SIFS
ACK
SIFS
FRAGMENT 2
(H C F POLL)
UPLINK
DOWNLINK
START T X O P
SIFS
FRAGMENT 1
TIMEOUT ACK
FRAGMENT 1
(H C F POLL)
UPLINK
DOWNLINK

FIG. 20

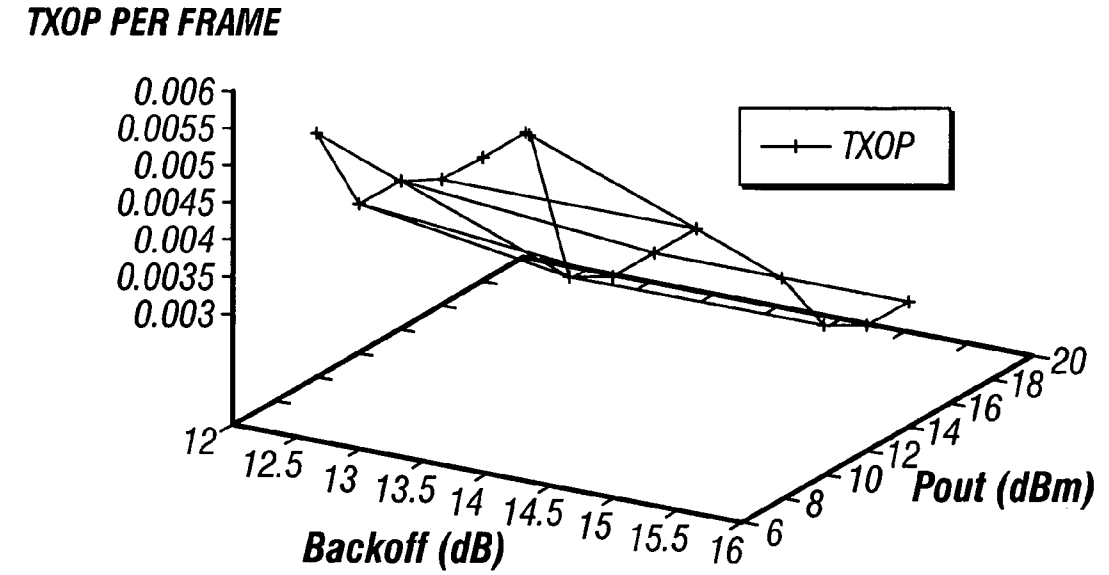


FIG. 21A

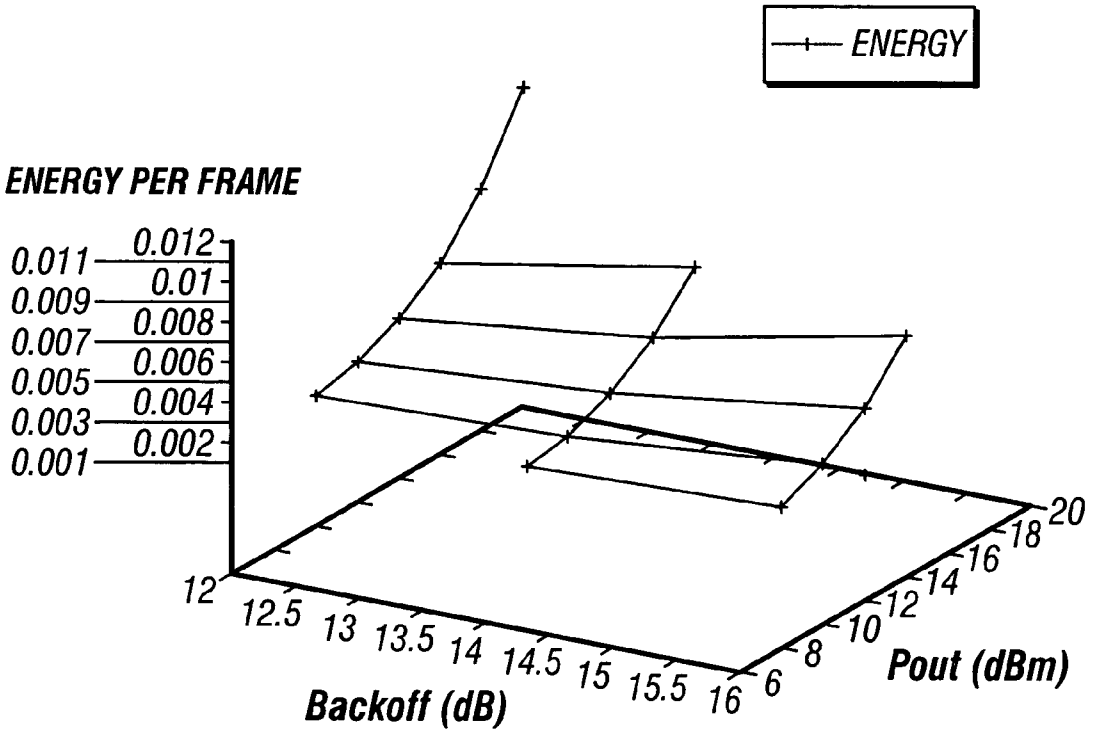


FIG. 21B

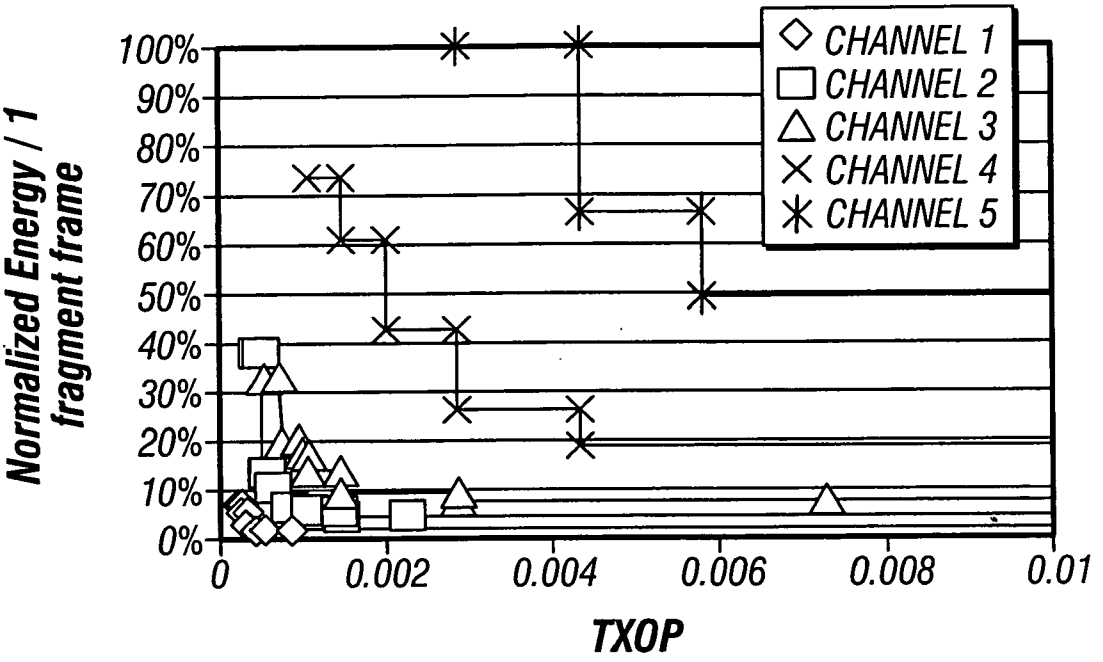


FIG. 22A

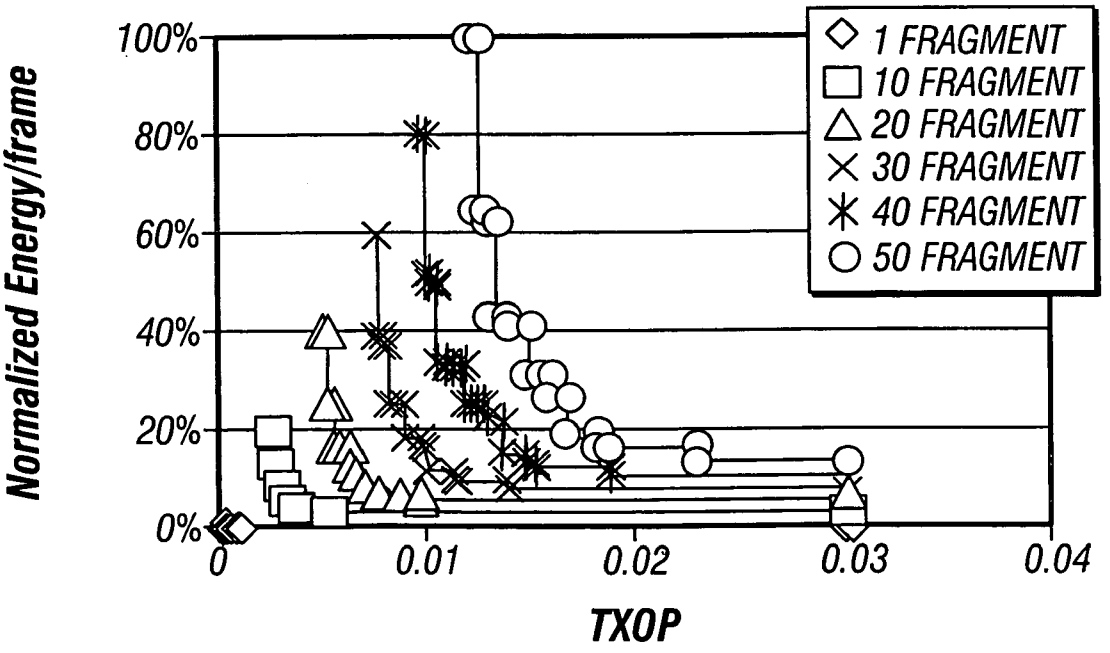


FIG. 22B

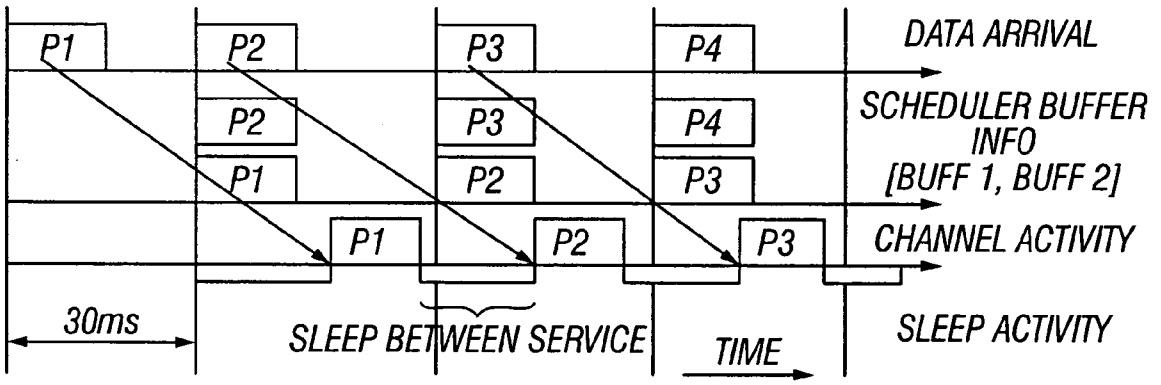

P1
P2
P3
P4
DATA ARRIVAL
P2
P3
P4
SCHEDULER BUFFER INFO [BUFF 1, BUFF 2]
P1
P2
P3
P1
P2
P3
CHANNEL ACTIVITY
30ms
SLEEP BETWEEN SERVICE
TIME
SLEEP ACTIVITY

FIG. 23

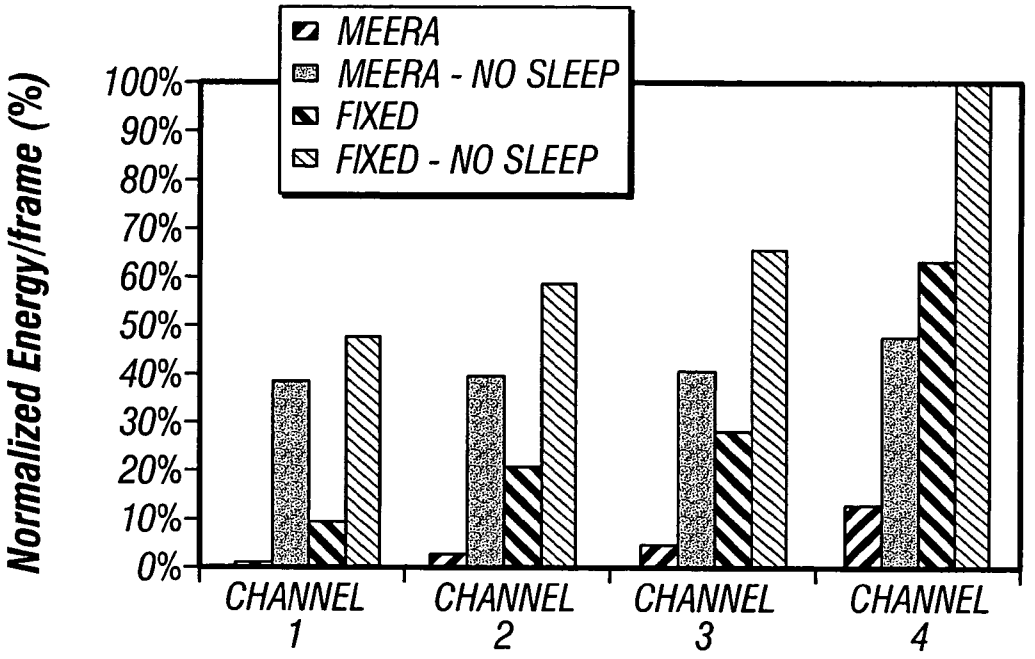


FIG. 24A

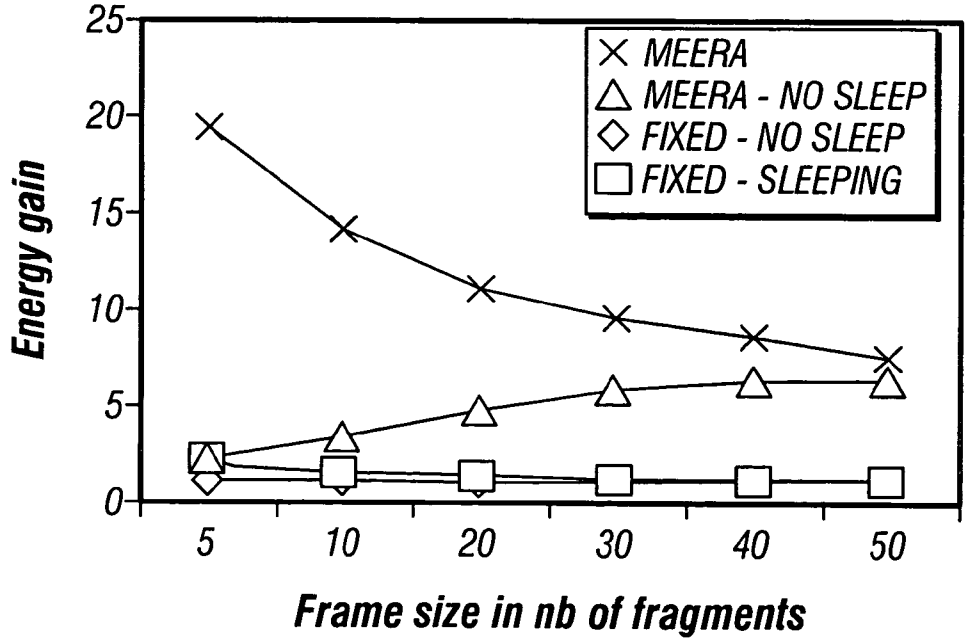


FIG. 24B

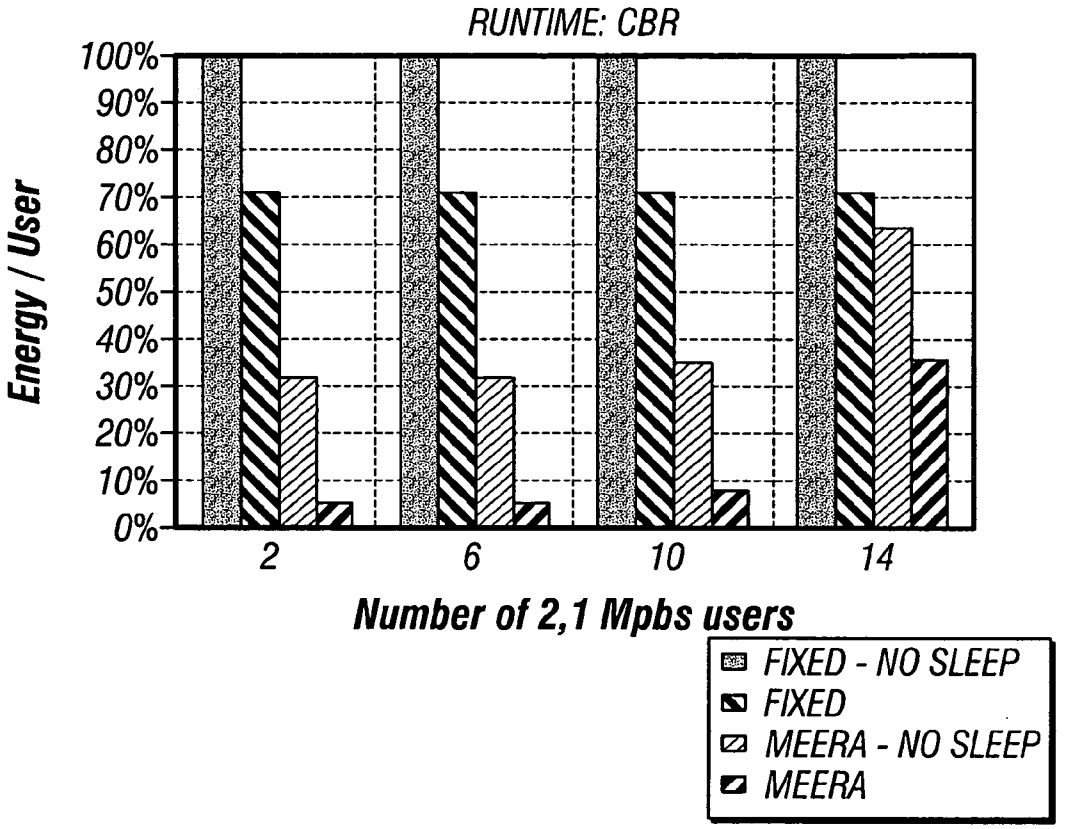


FIG. 25A

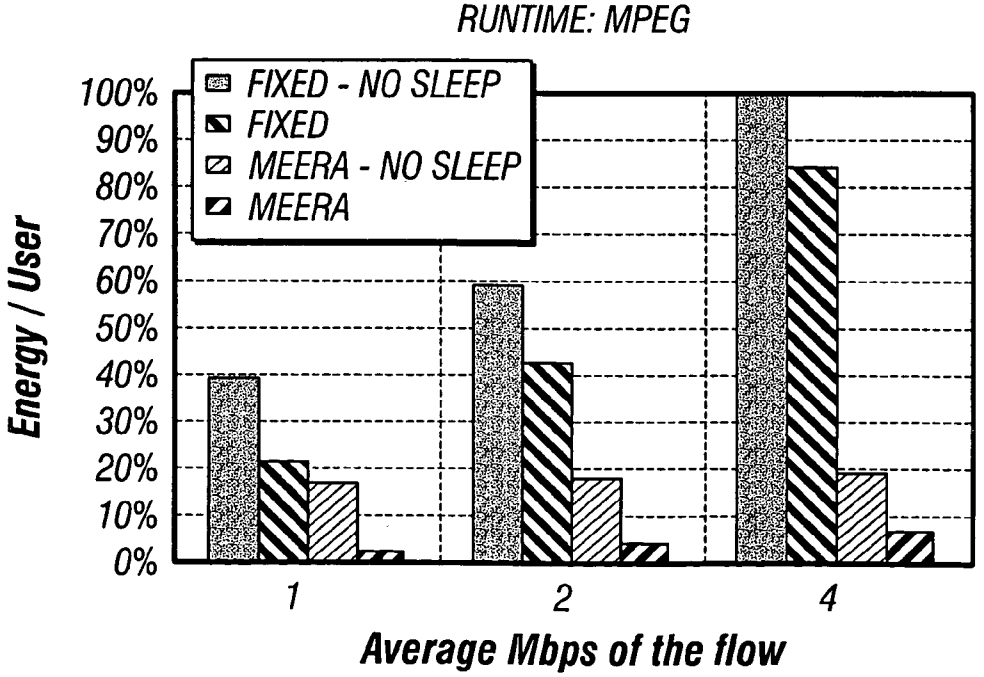


FIG. 25B

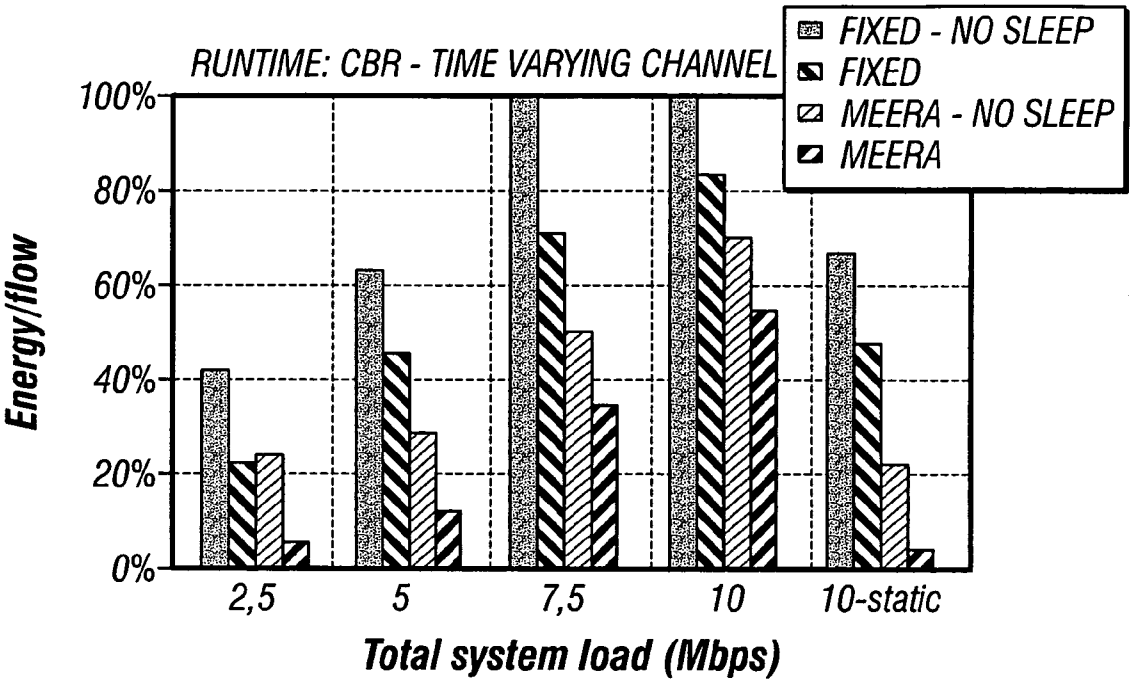


FIG. 26A

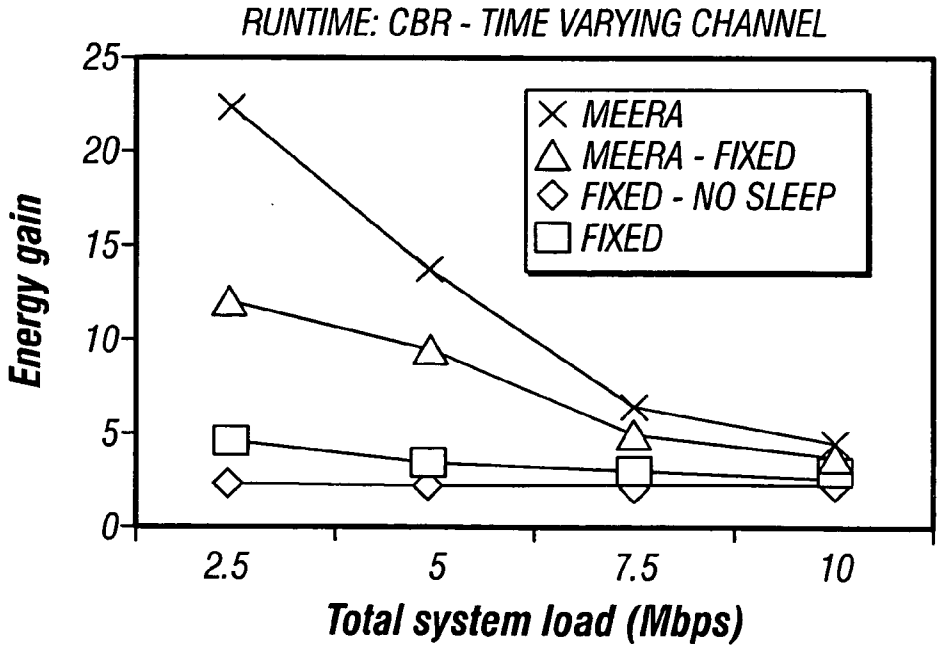


FIG. 26B

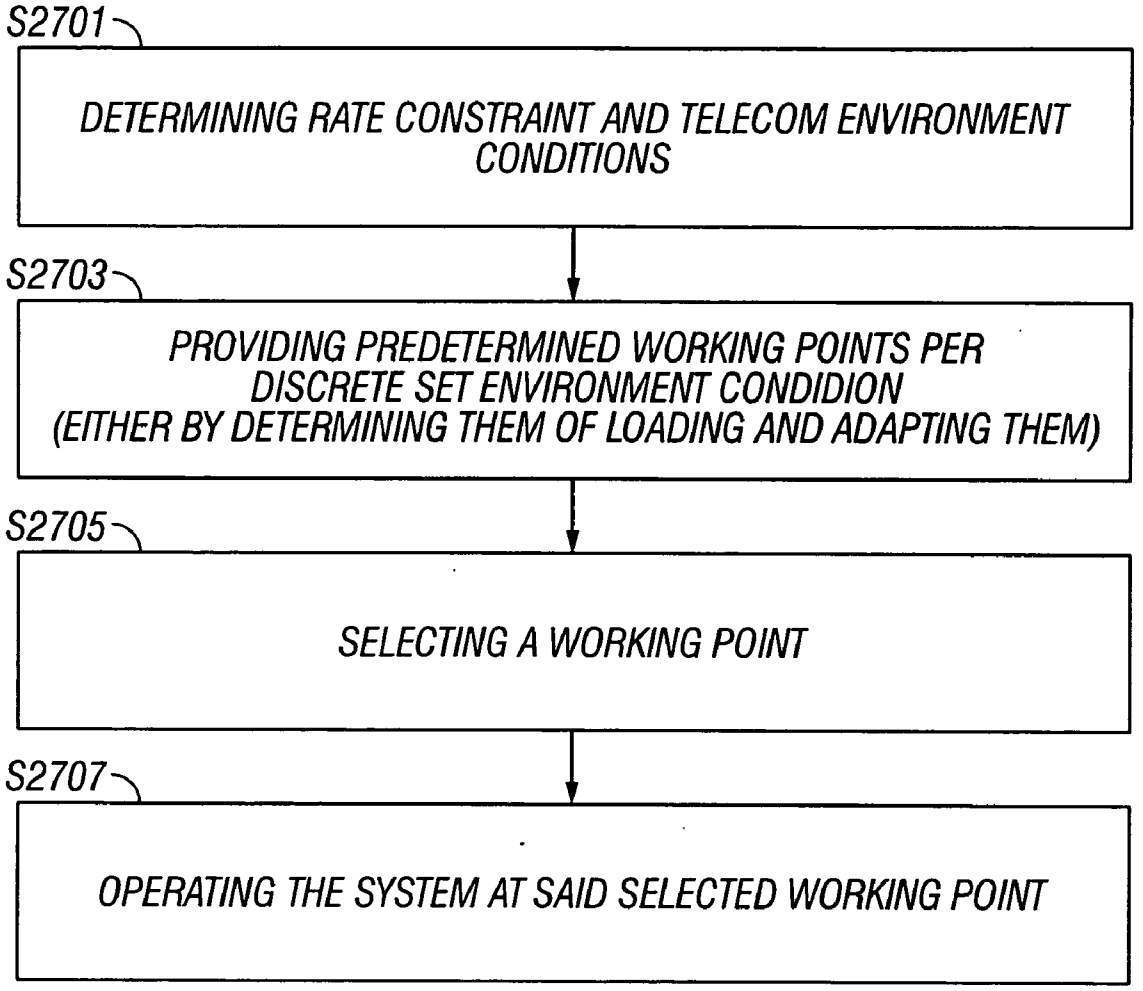
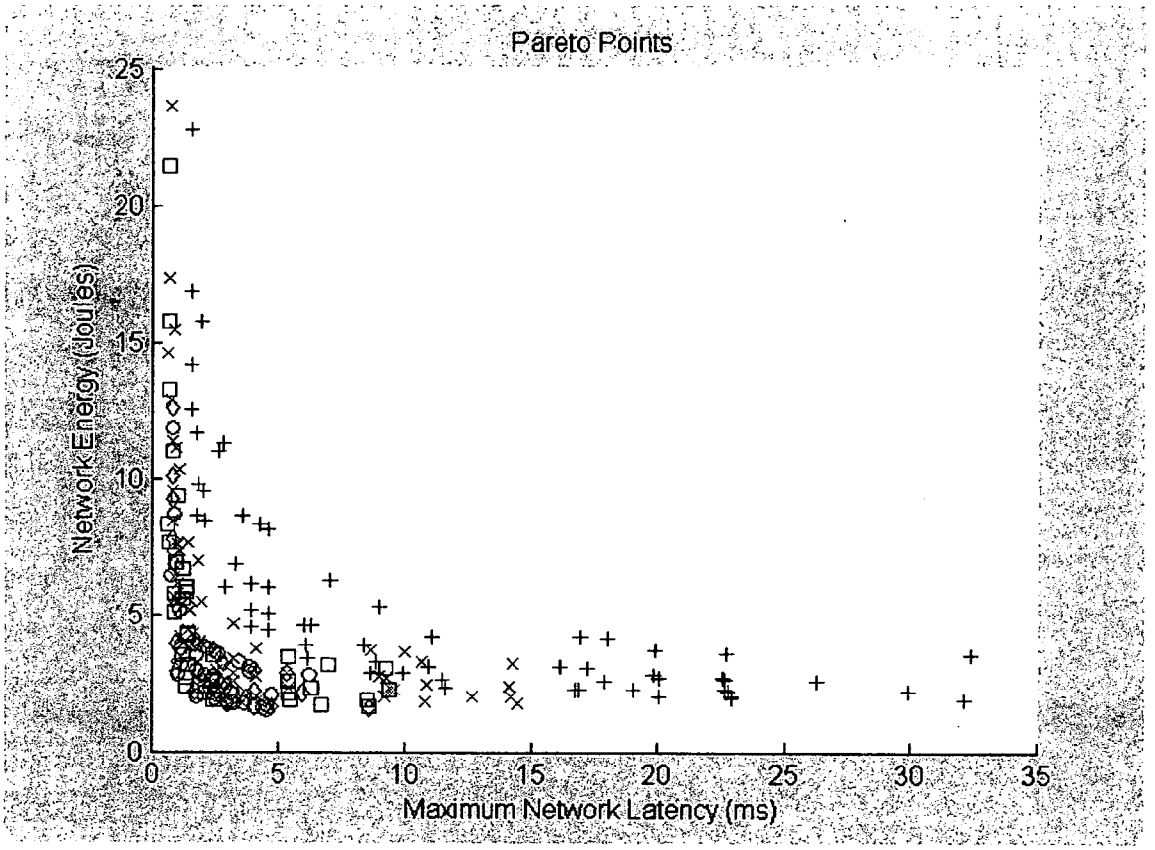
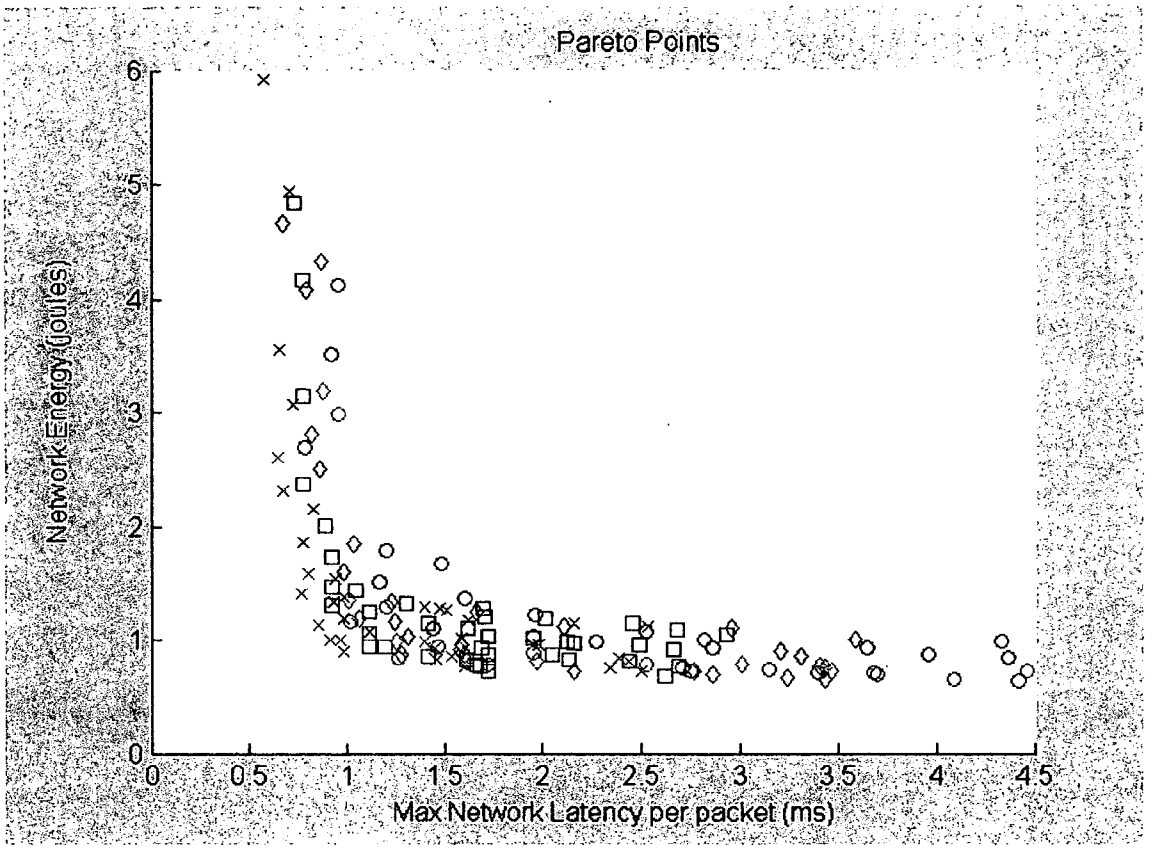

S2701
DETERMINING RATE CONSTRAINT AND TELECOM ENVIRONMENT CONDITIONS
S2703
PROVIDING PREDETERMINED WORKING POINTS PER DISCRETE SET ENVIRONMENT CONDIDION (EITHER BY DETERMINING THEM OF LOADING AND ADAPTING THEM)
S2705
SELECTING A WORKING POINT
S2707
OPERATING THE SYSTEM AT SAID SELECTED WORKING POINT

FIG. 27



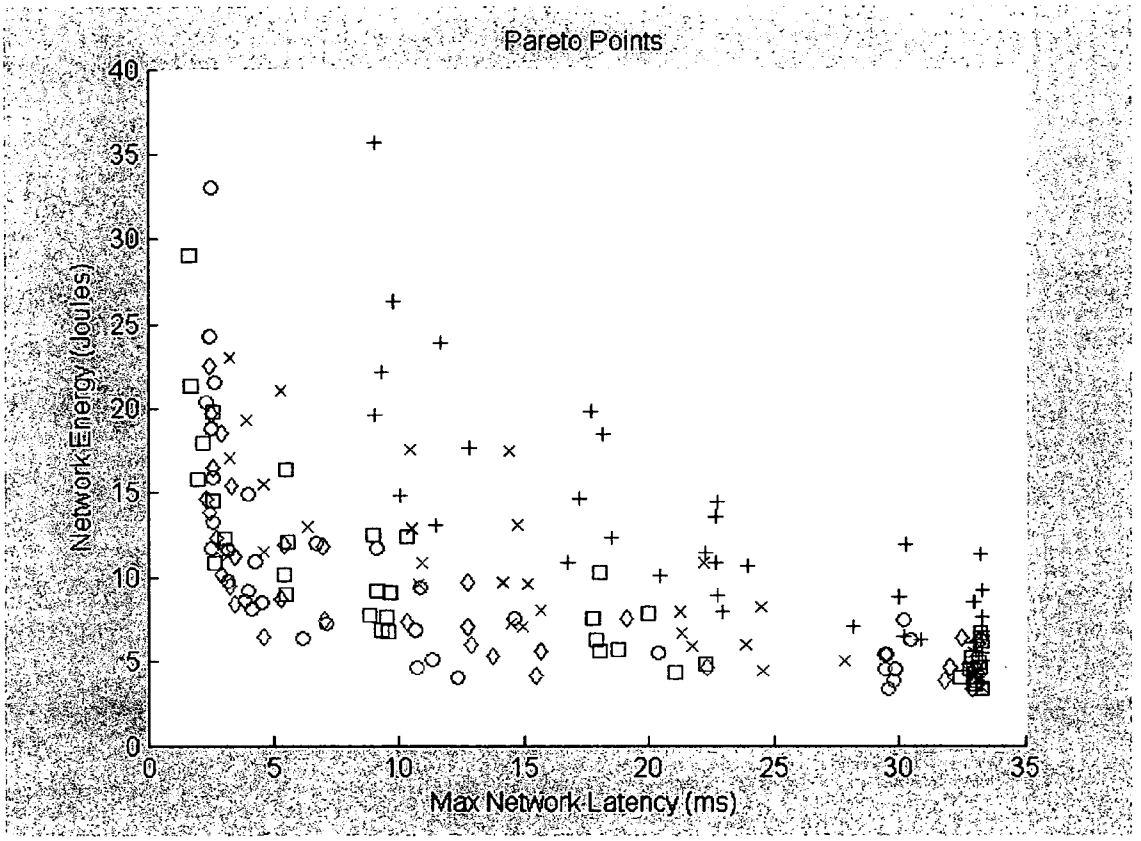
DCF mode, Mobile sequence, Channel 2, Network Energy and Latency

FIG. 28



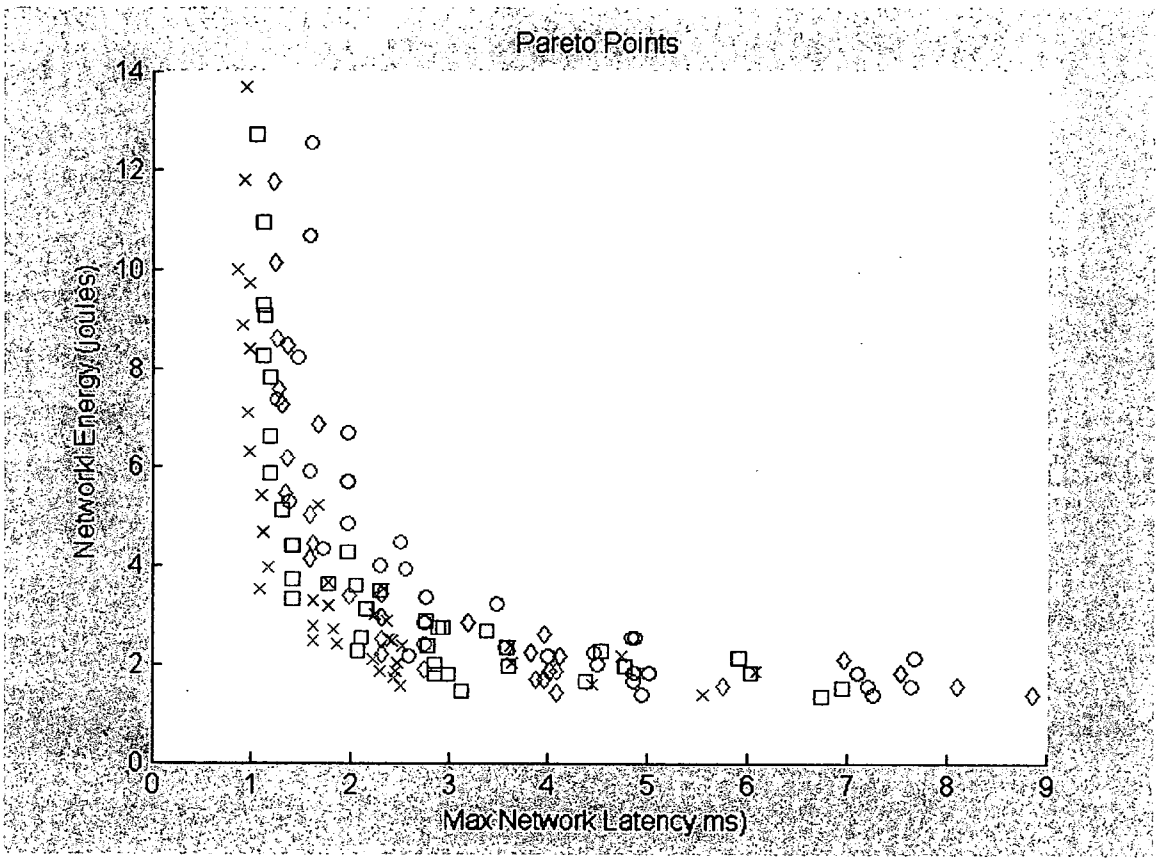
DCF mode, Foreman sequence, Channel 2, Network Energy and Latency

FIG. 29



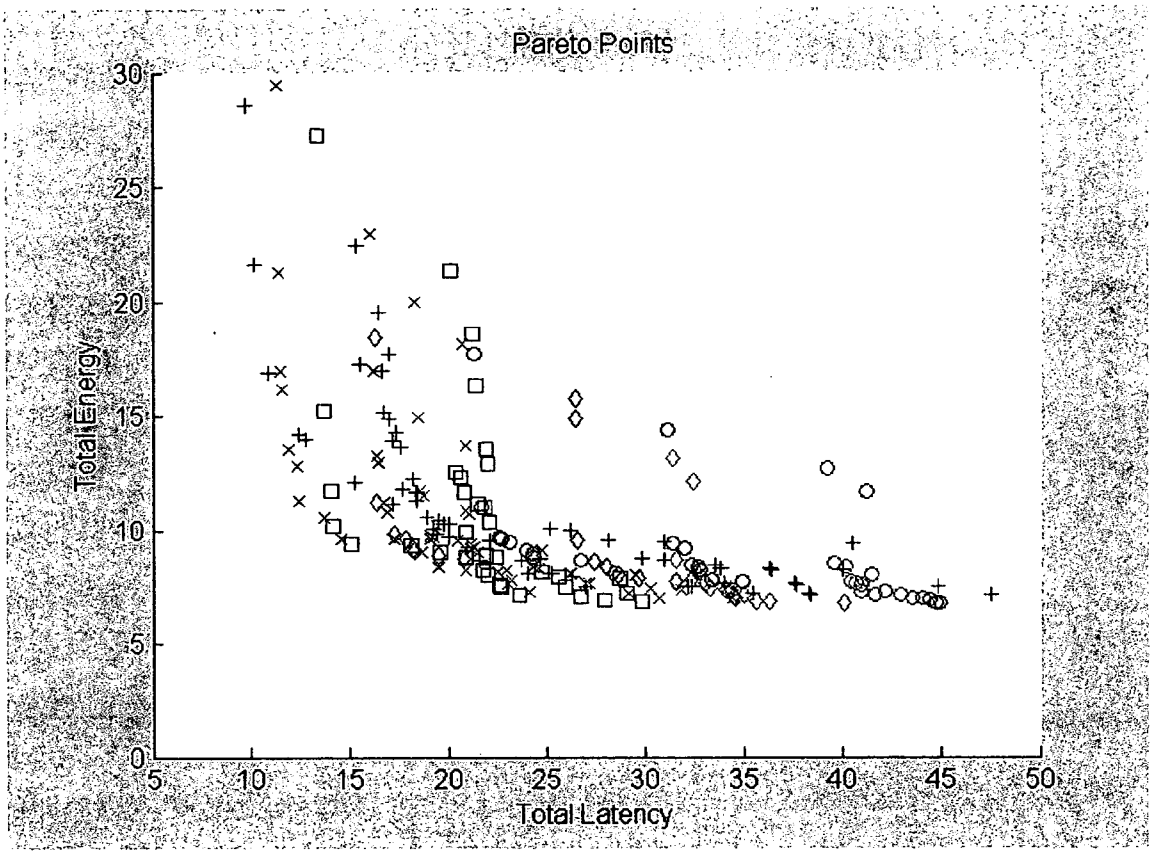
DCF mode, Mobile sequence, Channel 4, Network Energy and Latency

FIG. 30



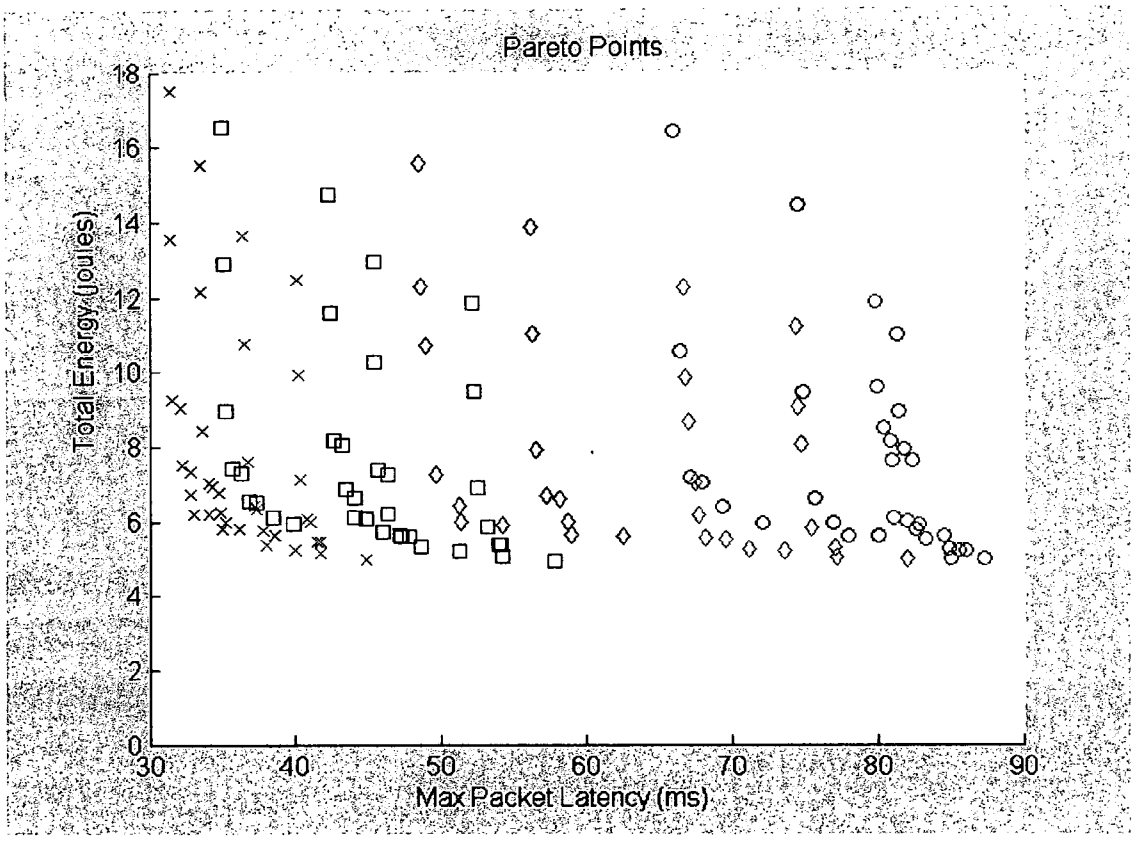
DCF mode, Foreman sequence, Channel 4, Network Energy and Latency

FIG. 31



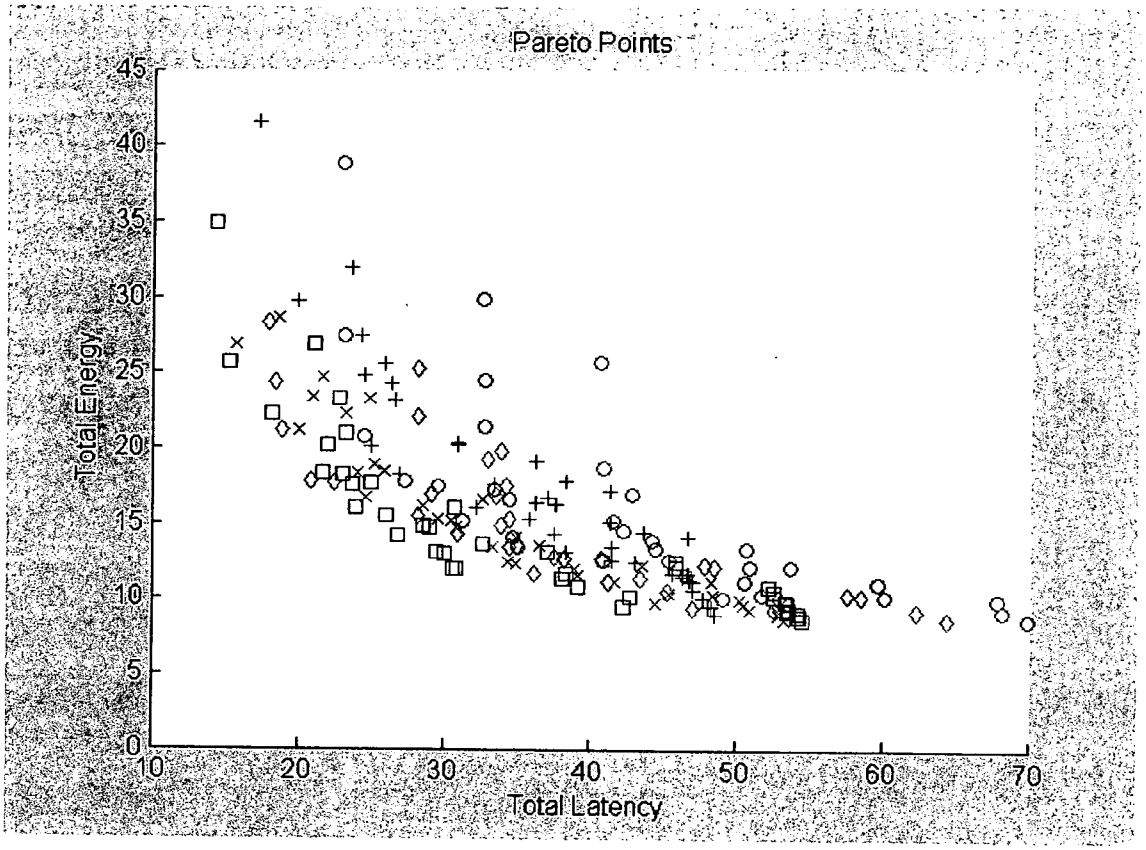
DCF mode, Mobile sequence, Channel 4, Network and Application Energy and Latency

FIG. 32



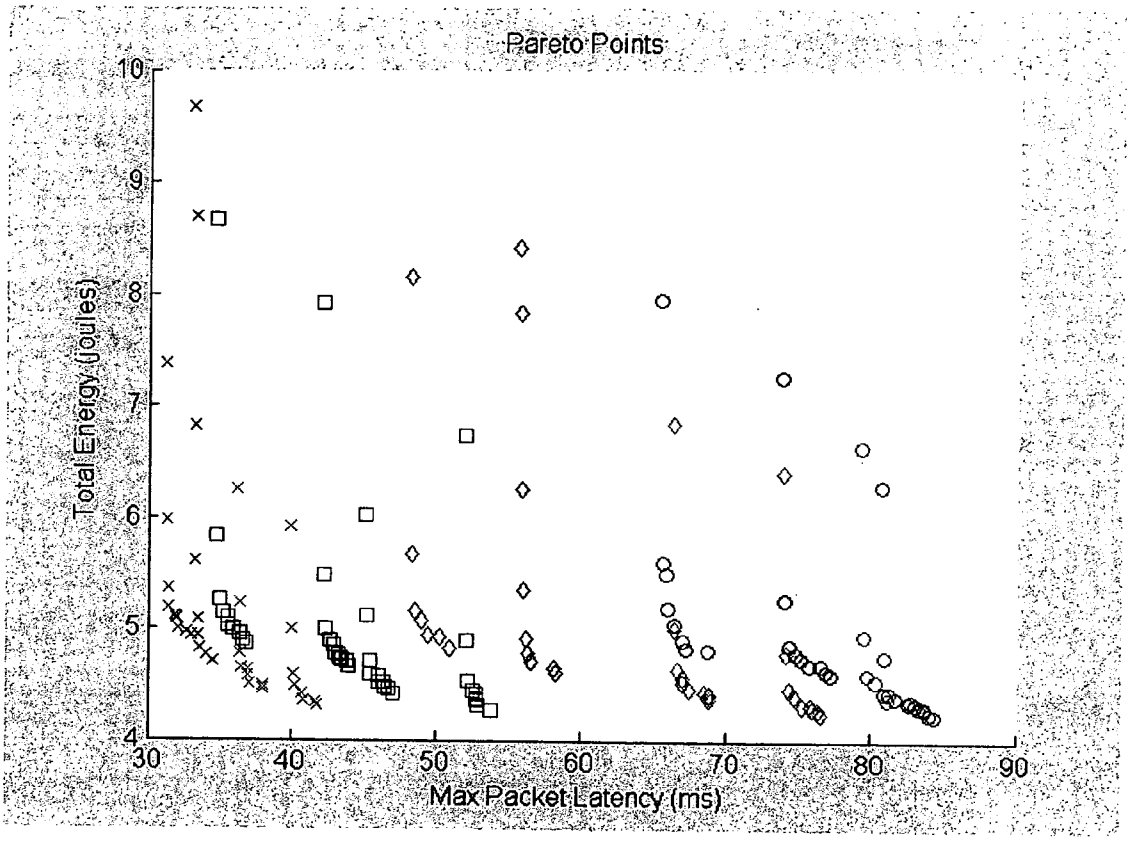
DCF mode, Foreman sequence, Channel 2, Network and Application Energy and Latency

FIG. 33



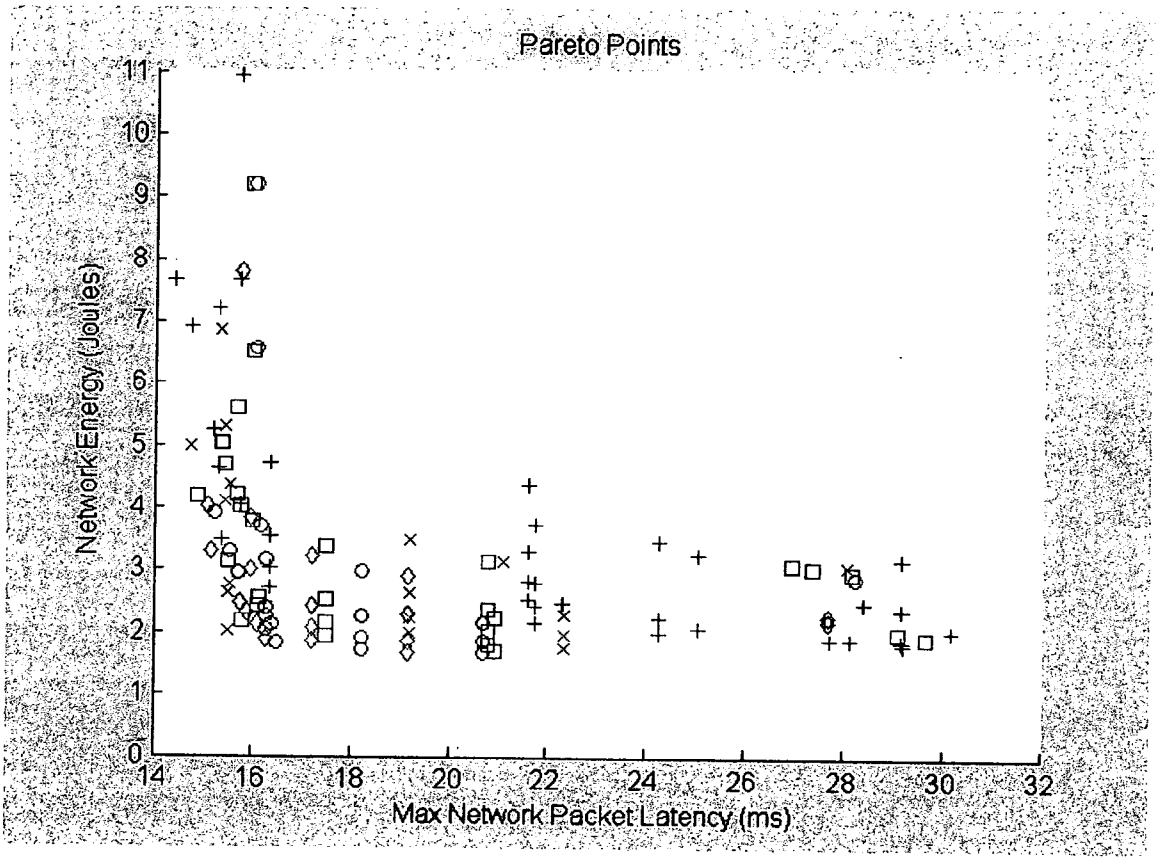
DCF mode, Mobile sequence, Channel 4, Network and Application Energy and Latency

FIG. 34



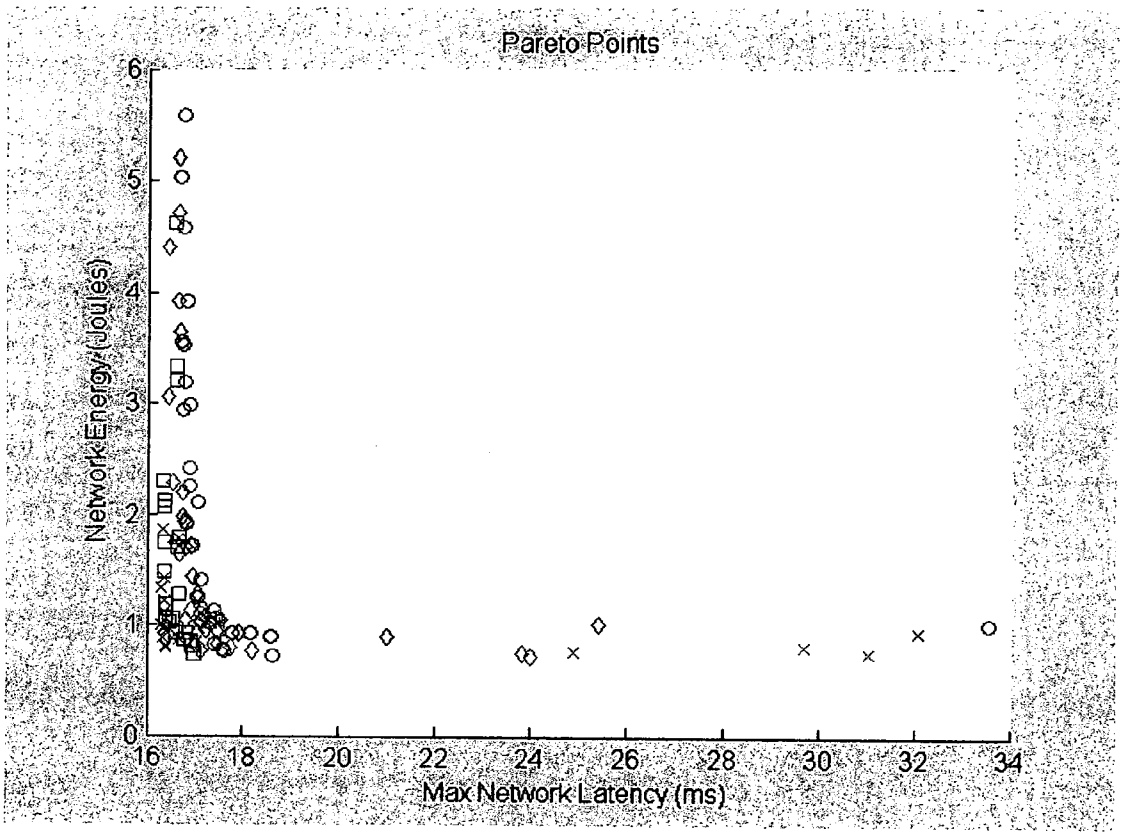
DCF mode, Foreman sequence, Channel 4, Network and Application Energy and Latency

FIG. 35



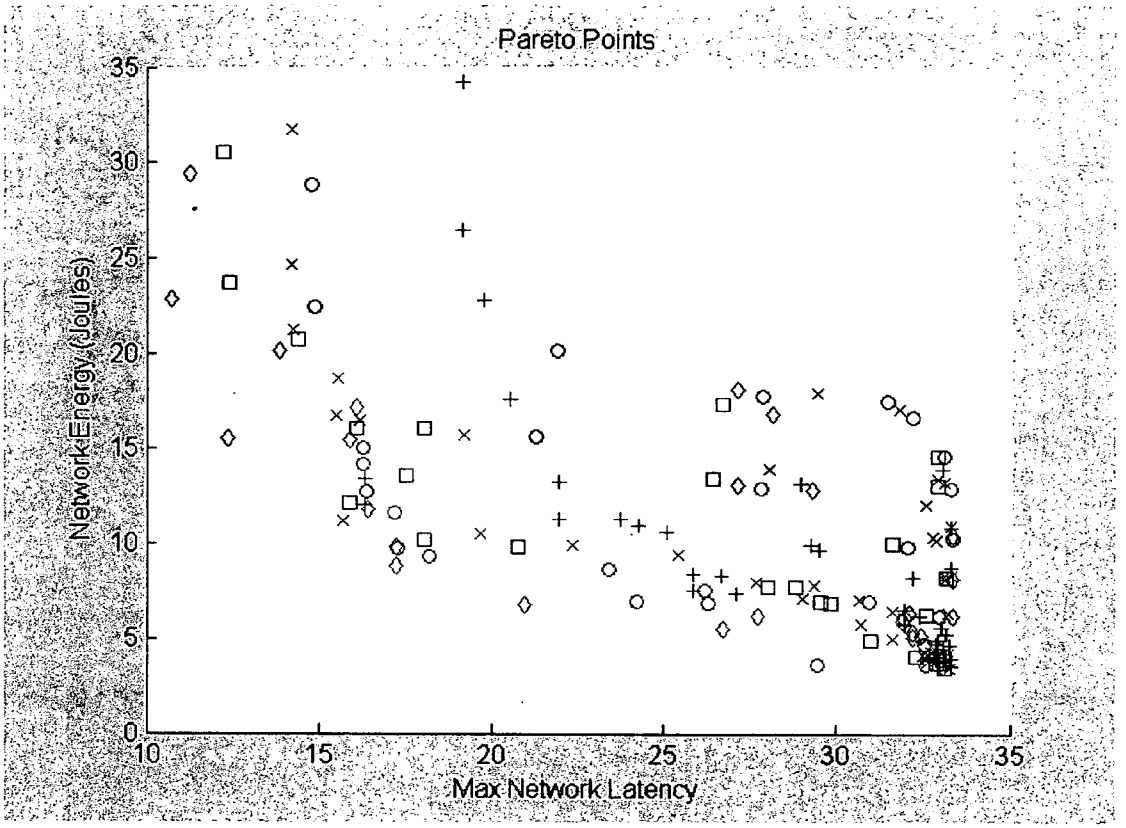
HCF mode with beacon of 17 ms, Mobile sequence, Channel 2, Network Energy and Latency

FIG. 36



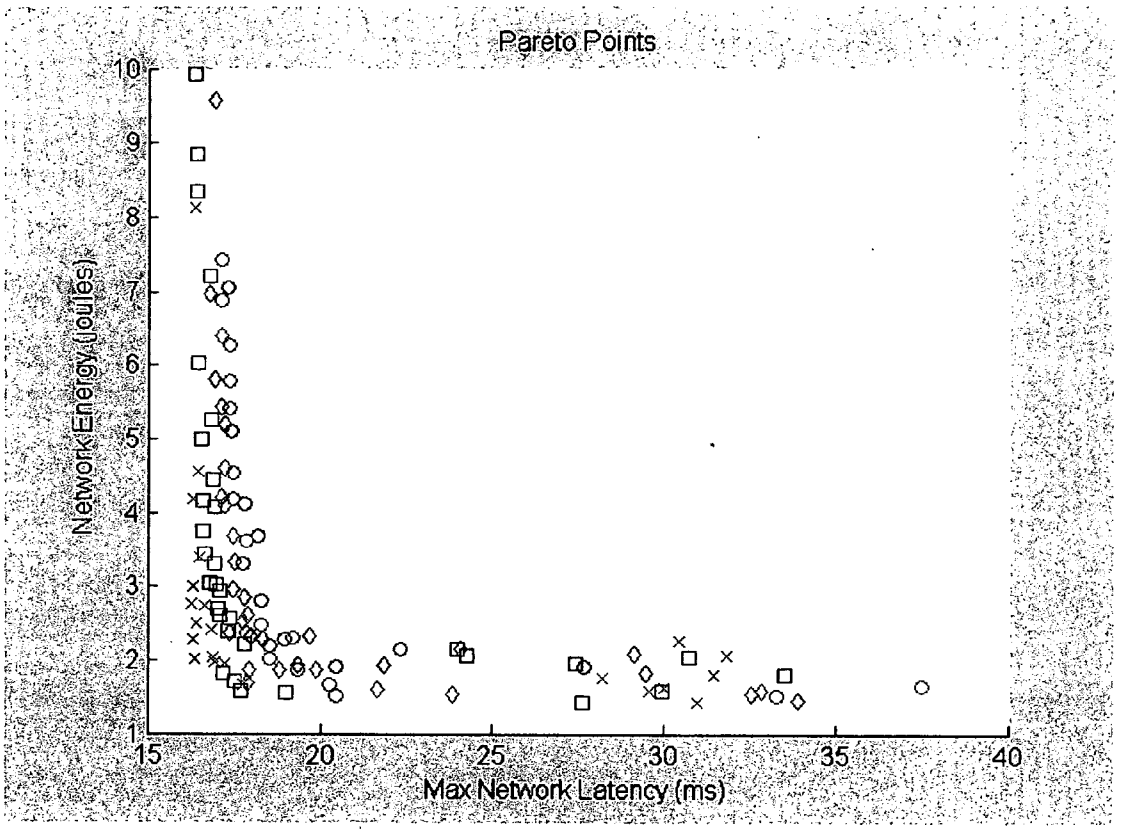
HCF mode with beacon of 17 ms, Foreman sequence, Channel 2, Network Energy and Latency

FIG. 37



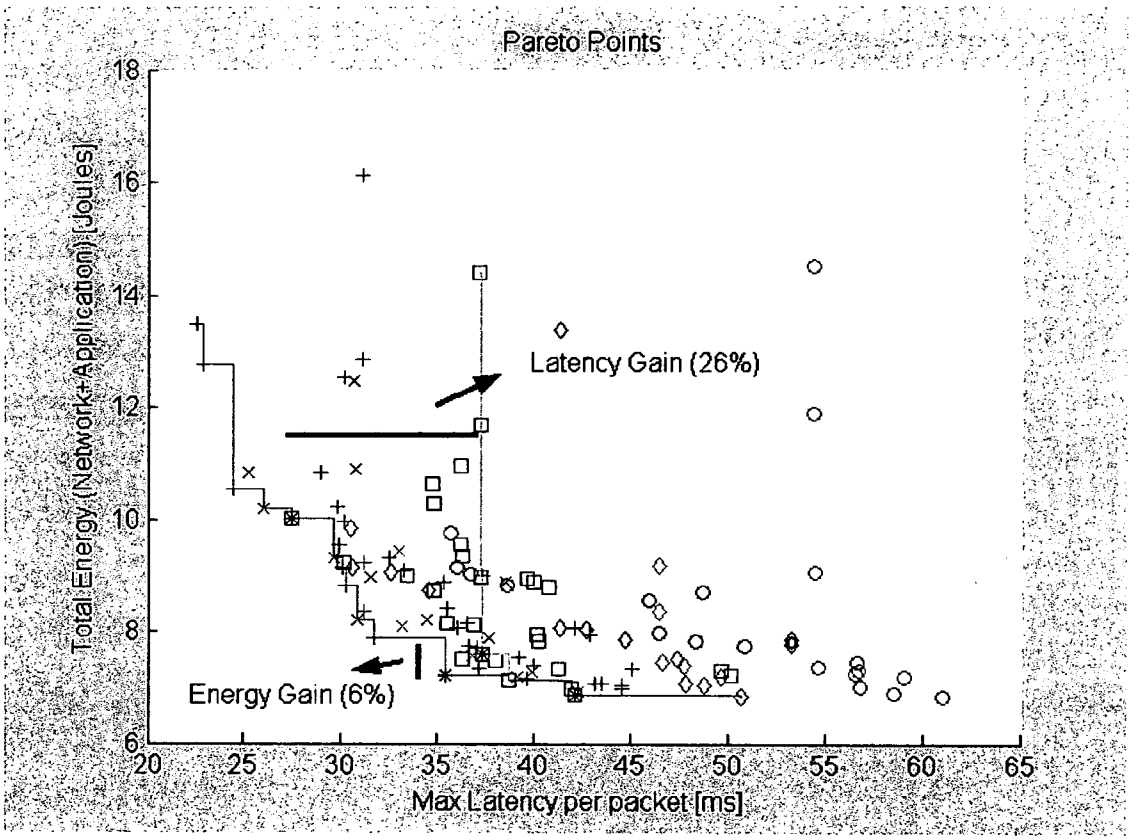
HCF mode with beacon of 17 ms, Mobile sequence, Channel 4, Network Energy and Latency

FIG. 38



HCF mode with beacon of 17 ms, Foreman sequence, Channel 4, Network Energy and Latency

FIG. 39



HCF mode with beacon of 17 ms, Mobile sequence, Channel 2, Network and Application Energy and Latency

FIG. 40

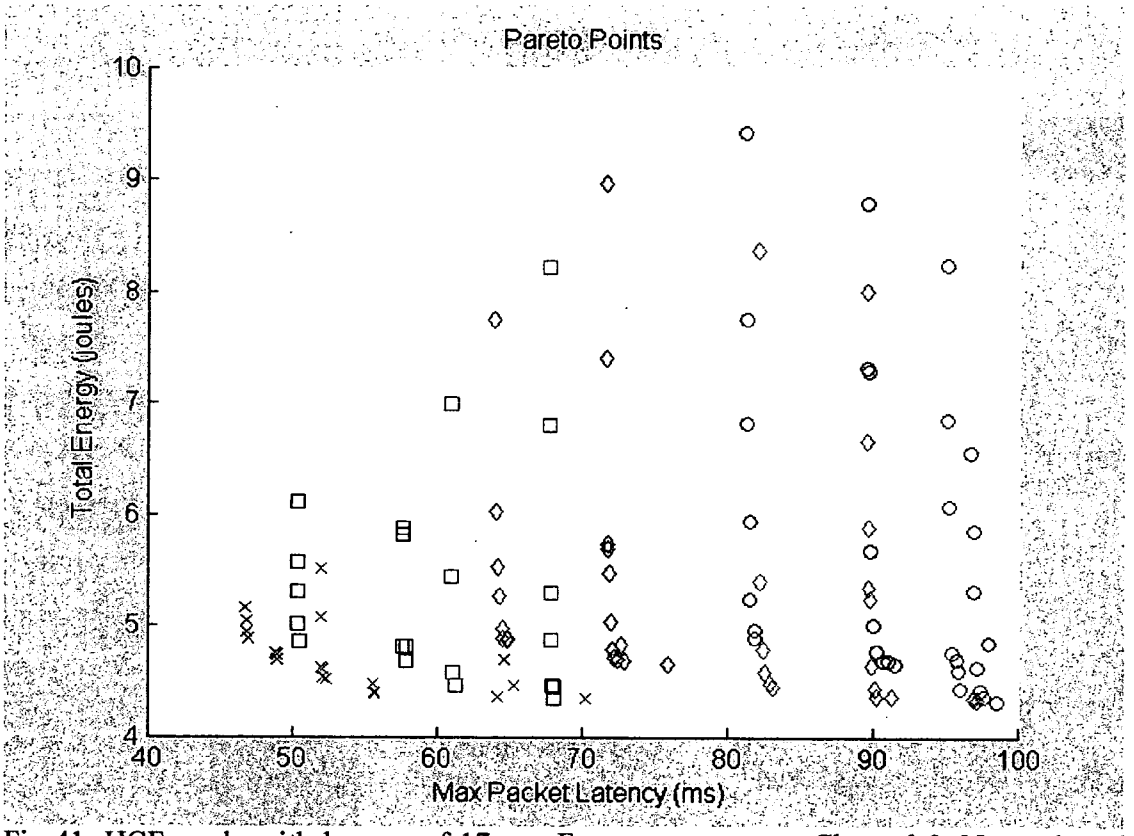
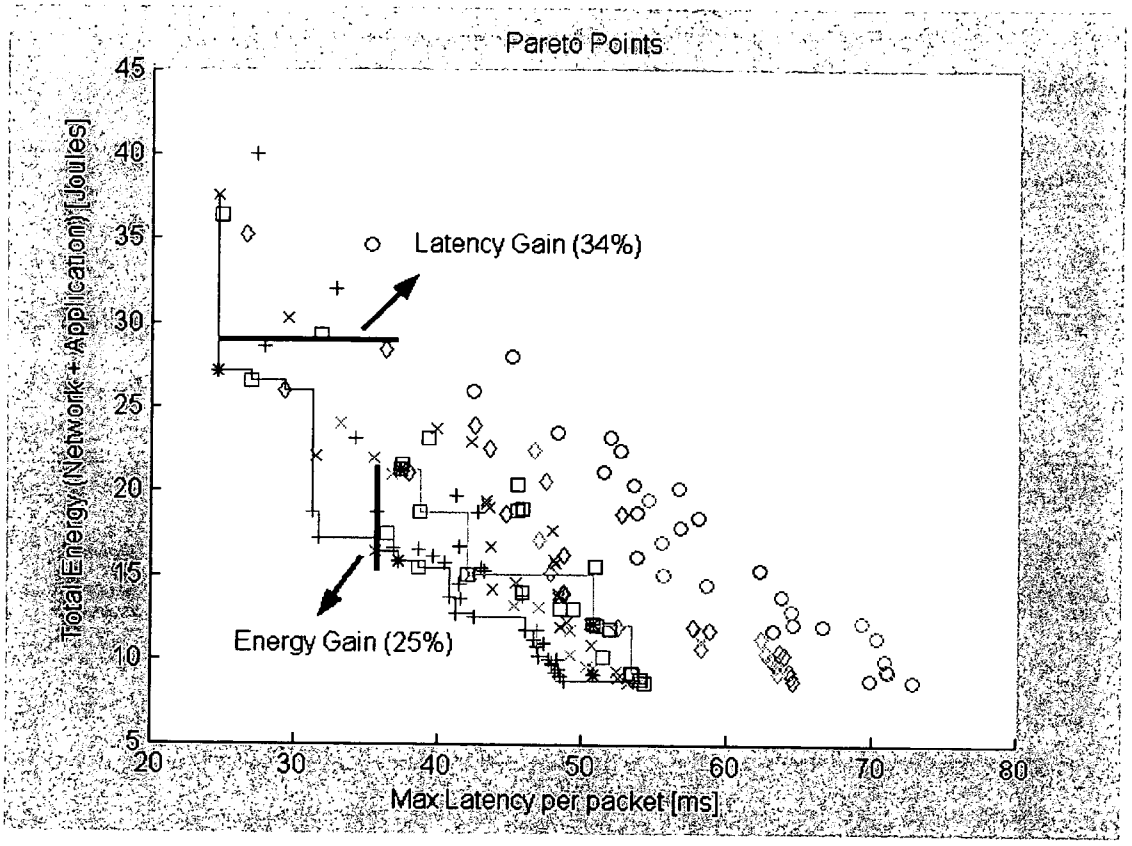


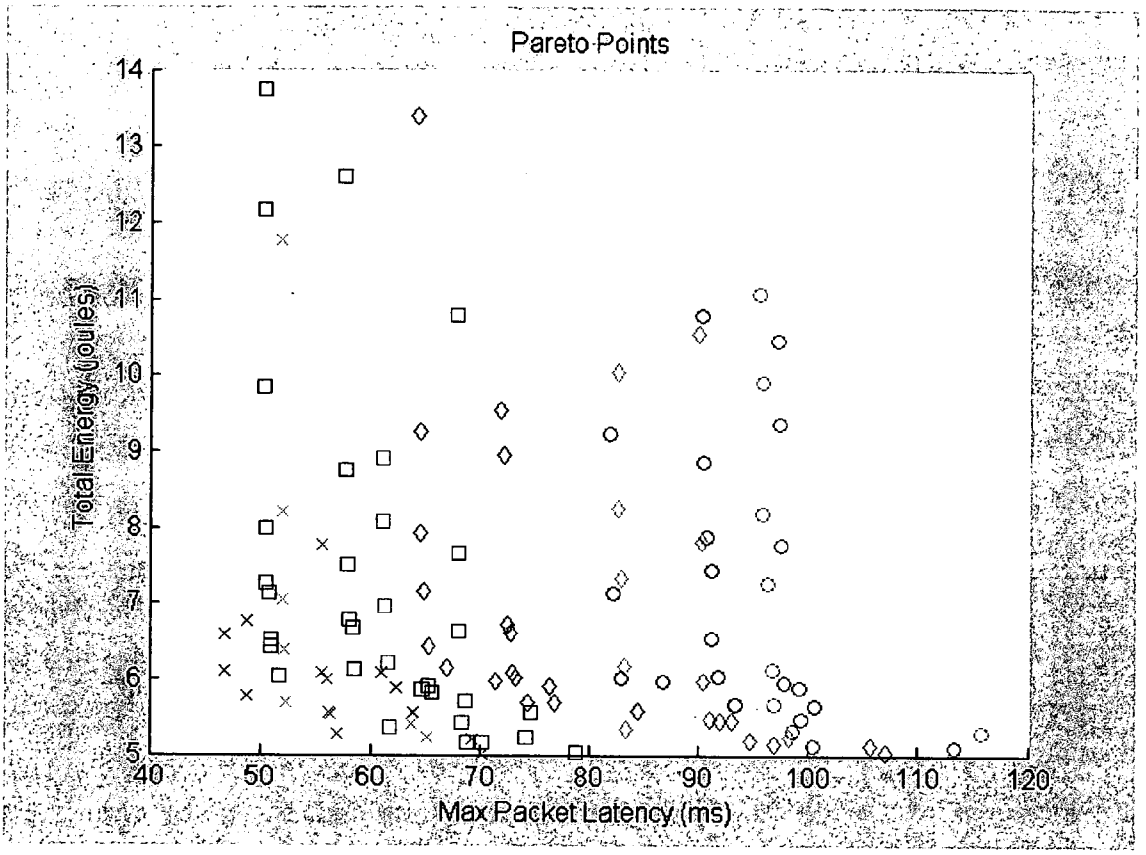
Fig 41. HCF mode with beacon of 17 ms, Foreman sequence, Channel 2, Network and Application Energy and Latency

FIG. 41



HCF mode with beacon of 17 ms, Mobile sequence, Channel 4, Network and Application Energy and Latency

FIG. 42



HCF mode with beacon of 17 ms, Foreman sequence, Channel 4, Network and Application Energy and Latency

FIG. 43

METHOD FOR OPERATING A COMBINED MULTIMEDIA -TELECOM SYSTEM

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation in part of U.S. application Ser. No. 10/922,371, titled "Method for Operating a Telecom System," filed Aug. 20, 2004, which is hereby incorporated by reference, in its entirety. This application also claims priority to U.S. Provisional Application No. 60/617,897, titled "Method for Operating a Combined Multimedia-Telecom System," filed Oct. 12, 2004, which is hereby incorporated by reference, in its entirety.

BACKGROUND OF THE INVENTION

[0002] 1. Field of the Invention

[0003] The present invention is related to a method for operating a combined multimedia telecom system, more in particular a combined video encoding-wireless system, and devices suited therefore.

[0004] 2. Description of the Related Technology

[0005] By discussion of technologies and references in this section, Applicants do not admit that the references are prior art of the invention disclosed in this application.

[0006] The current demand in increasing data rate and quality of service in advanced wireless communications has to cope with an energy budget severely constrained by autonomy and ubiquity requirements. Trading off performance and energy consumption deserves the highest attention to enable the 'anything, anywhere, anytime' paradigm.

[0007] The following observations show the need to integrate the energy-efficient approaches across layers. First, state-of-the-art wireless systems devices are built to operate at only a fixed set of operating points and assume the worst-case conditions at all times. Irrespective of the link utilization, the highest feasible PHY rate is always used and the power amplifier operates at the maximum transmit power. Indeed, when using non-scalable transceivers, this highest feasible rate results in the smallest duty cycle for the power amplifier. Compared to scalable systems, this results in excessive energy consumption for average channel conditions and average link utilizations. However, recent energy-efficient wireless system designs focus on energy-efficient VLSI implementations and adaptive physical layer algorithms to adjust the modulation, code rate or transmission power. For these schemes to be practical, they need to be aware of the instantaneous user requirements.

[0008] Further, to realize sizeable energy savings, systems need to shutdown the components when inactive. This is achieved only by tightly coupling the MAC to be able to communicate traffic requirements of a single user and schedule shutdown intervals.

[0009] In the case of a multi-user wireless communication system, there exist complex trade-offs between the adaptive physical layer schemes and the requirements of multiple users. For example, lowering the rate of one user affects the available time for the second delay sensitive user. This forces the second user to increase its rate, consume more energy and potentially suffer from a higher bit error rate.

[0010] However, the traditional approaches, including most of the state-of-the-art cross-layer optimization frameworks, do not yet enable a meaningful trade-off between performance and energy consumption. Indeed, most of them solve problems in an ad hoc way, focusing on the interaction between adjacent layers and do not raise the scope to the user level. Indeed, the real performance metrics are those quantifying the quality of the service provided by the complete communication stack to the application, while the only effective energy consciousness indicator is the actual energy that is drained from the battery. Both depend jointly on the propagation aspects, the physical layer, the complete protocol stack, the application itself and, more problematically, also on their implementation aspects. This spans far more than the scope of traditional system optimization approaches. Furthermore, the traditional 'optimization paradigm' itself, namely finding a unique optimal communication system configuration representing the best trade-off between performance and cost, becomes inappropriate when the dynamics in the wireless environment and user requirements are considered. More specifically, because of this dynamics, no unique optimal working point exists. The ultimate energy-efficient system would have to adapt permanently its characteristic, given the environment constraints, to provide the performance exactly required by the user with the minimum energy.

[0011] To achieve this goal flexible systems must be specified having so-called configuration knobs that can be set at run-time to steer jointly the performance and energy consumption. The higher the flexibility, i.e. the number of configuration knobs across all layers, the higher the potential gain due to a better match between the system behavior, the environment and the real user requirements. However, a penalty exists due to the required shift, at least partially, of the optimization process to run-time. This is very challenging due to the combinatorial character of the problem (the number of candidate configurations rises exponentially with the number of controlled knobs).

[0012] Recently, joint transmission power and rate control has been considered to reduce system power (see D. Qiao et al., '*Energy Efficient PCF Operation of IEEE*802.11*a WLAN with Transmit Power control*', Elsevier Computer Networks (ComNet), vol. 42, no. 1, pp. 39-54, May 2003 and '*MiSer: An Optimal Low-Energy transmission Strategy for IEEE* 802.11*a/h', Proc. ACM MobiCom* '03, San Diego, September 2003). This approach can be seen as the application to wireless system design of the 'power aware' design paradigm proposed by Sinha et al. ('*Energy Scalable System Design', Trans. on VLSI Systems*, April 2002, pp. 135-145). Given the fact that transmitting at a lower rate requires less power, the 'lazy scheduling' principle has been proposed (see '*Adaptive Transmission for Energy Efficiency in Wireless Data Networks*', E. Uysal-Biyikoglu, Ph.D. Thesis, Stanford, June 2003): based on a look-ahead of the link utilization (i.e. packet arrival at the transmitter), the minimum average rate to satisfy the current traffic requirements is considered and transmit rate and power are set in function of the channel state in order to achieve this average during the next look-ahead window.

[0013] In '*Energy-aware Wireless Communications*' (C. Schurgers, Ph.D. thesis, University of California, Los Angeles, 2002) the concept of energy-aware radio-management is developed. It proposes simple models to capture the energy

consumption of radio systems that are used to derive some energy-efficient algorithms to select the modulation order, the code rate and to schedule the packets. This dynamic modulation and code scaling is proposed as a practical way to implement lazy scheduling. It also discusses the energy trade-off between transmission rates and shutting off the system. Operating regions are derived when a transceiver may sleep or use transmission scaling for time-invariant and time-varying channels. However, the general solutions to transparently exploit the energy-performance scalability at run-time are limited to a few (2-3) system level knobs. The energy-scalability of a system can be defined as the range in which the energy consumption can vary when the performance requirements—e.g. the user data rate—or the external constraints—e.g. the propagation conditions—vary from worst to best case.

[0014] In '*Practical Lazy Scheduling in Sensor Networks*', R. Rao et al, Proc ACM Sensor Networks Conf, Los Angeles, November 2003 a CSMA/CA MAC protocol based on the lazy scheduling idea is derived.

[0015] From a theoretical point of view, the 'lazy scheduling' concept is attractive. E.g. radio link control based on 'lazy scheduling' looks to be a promising technique for WLAN power management. However, it has been analyzed exclusively from the viewpoint of physical, MAC and dynamic link control (DLC) layer. Yet, effective power management in radio communication requires considering the complete protocol stack and its cross-layer interactions.

SUMMARY OF THE INVENTION

[0016] Certain embodiments provide a method for operating a telecom system, more in particular a wireless system, with a globally optimized power consumption for a given quality of service. Other embodiments further aim to provide devices suited to the method.

[0017] One embodiment includes a method of managing the operation of a system including a processing subsystem configured to run a multimedia application and a telecommunication subsystem. The method includes determining telecom environment conditions, and selecting a working point from a plurality of predetermined working points. The selecting is based at least in part on the determined environmental conditions. The method also includes setting control parameters in the multimedia application and/or the telecommunication subsystem to configure the system to operate at the selected working point, and operating the system at the selected working point.

[0018] Another embodiment includes a computer readable medium storing a program configured to execute a method of managing the operation of a system including a processing subsystem configured to run a multimedia application and a telecommunication subsystem. The method includes determining telecom environment conditions, and selecting a working point from a plurality of predetermined working points. The selecting is based at least in part on the determined environmental conditions. The method also includes setting control parameters in the multimedia application and/or the telecommunication subsystem to configure the system to operate at the selected working point, and operating the system at the selected working point.

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] The patent or application file contains at least one drawing executed in color. Copies of this paten or patent application publication with color drawing(s) will be provided by the Office upon request and payment of the necessary fee.

[0020] **FIG. 1** represents the energy consumption breakdown in typical wireless transceivers.

[0021] **FIG. 2A** represents DC power consumption of a Class A power amplifier (PA).

[0022] **FIG. 2B** represents adaptation of the PA gain compression characteristic.

[0023] **FIG. 3** represents the physical layer simulation chain used to derive the PER performance model.

[0024] **FIG. 4A** represents block error rate performance of 8 states of the fading channel realizations.

[0025] **FIG. 4B** is a histogram of the state probabilities of the 8 states of the fading channel realizations.

[0026] **FIGS. 5A-5G** are the optimal energy performance trade-off curves for a 802.11a WLAN link at 10 m, for standard transceivers (dashed line) and energy-scalable transceiver (plain line).

[0027] **FIG. 6** represents a variable rate transmission system model.

[0028] **FIG. 7** represents the average power versus average delay trade-off for the optimal (hypothetical) policy (square), the original lazy scheduling (circle) and the proposed proportional control policy (triangle).

[0029] **FIG. 8** represents the average power vs. average delay for different link load and different values of K. The delay axis is normalized by the control period.

[0030] **FIGS. 9A-9K** represent the average delay as a function of the control parameter K in various input bit rates. Up to 80% utilization, the average delay is inversely proportional to K. At higher utilization, the proportionality is lost, the queuing behavior becoming non-linear due to congestion of the channel.

[0031] **FIG. 10** represents a wireless point-to-point file transfer scenario.

[0032] **FIG. 11** represents the packet drop probability as a function of K.

[0033] **FIG. 12** represents TCP throughput per connection (i.e. one out of the 10 connections sharing the link) as a function of the control parameter K.

[0034] **FIG. 13** represents the energy per byte versus TCP throughput trade-off achievable by tuning the link scheduling policy.

[0035] **FIG. 14** represents a distributed stack-wide power management.

[0036] **FIG. 15** represents Pareto curves obtained after calibration for a path loss of 80 dB.

[0037] **FIG. 16** represents the average power versus average goodput trade-off achieved by the proposed radio link control strategy (circle) compared to the traditional goodput maximizing strategy. The plain line corresponds to the optimal selection.

[0038] **FIG. 17** represents a centrally controlled LAN/PAN topology illustrating uplink and peer-to-peer communication.

[0039] **FIG. 18** represents the determination of a pruned $C_i(R_i)$ mapping to be used at runtime, starting from a cost and resource profile.

[0040] **FIG. 19A** represents the performance across different channel states.

[0041] **FIG. 19B** is a channel state histogram of a probability of occurrence based on the channel realizations database.

[0042] **FIG. 20** represents the timing of successful and failed uplink frame transmission under the HCF and the scheme.

[0043] **FIGS. 21A and 21B** represent the mapping for the PA output power and back-off control dimension for a fixed setting of the modulation and code rate control dimensions.

[0044] **FIG. 22A** represents $C_i(R_i)$ curves in different channel states for 1 fragment.

[0045] **FIG. 22B** represents $C_i(R_i)$ curves for different frame sizes in channel state 1.

[0046] **FIG. 23** represents the two-frame buffering.

[0047] **FIG. 24A** represents expected energy consumption across different channel states for 1 fragment frame size.

[0048] **FIG. 24B** represents relative energy consumption by sleeping and scaling for different system loads in the best channel state.

[0049] **FIG. 25A** represents energy consumption per flow as a function of the aggregate system load for CBR traffic.

[0050] **FIG. 25B** represents energy consumption per flow as a function of mean per-flow data rate for MPEG traffic.

[0051] **FIG. 26A** represents energy consumption for CBR traffic over a time-variant channel as function of aggregate system load

[0052] **FIG. 26B** represents energy gains for sleeping versus scaling for time-variant channel.

[0053] **FIG. 27** is a flowchart of an embodiment for managing the operation of a telecom system.

[0054] FIGS. **28** to **43** show Pareto points obtained in the Energy-Latency space for DCF and HCF transmission modes.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

[0055] The following detailed description will explain various features of various embodiments in detail. The invention can be embodied in numerous ways as defined and covered by the claims. In order to clarify certain parts of this description some concepts will be illustrated with the practical example of an OFDM-based WLAN system. Some embodiments may be practiced within a QoS management system as described in U.S. Ser. No. 10/254,044, titled "Method for Operating a Real-Time Multimedia Terminal and a QoS Manner," filed Sep. 24, 2002 which is incorporated herein by reference in its entirety. Embodiments herein also relate to EP 1467294, titled "Design Method for Electronic Systems Using Library of Hardware Components with Performance Parameters and Cost Functions," filed Mar. 25, 2004, U.S. Ser. No. 10/252,688, titled "Method and Device for Block-Based Conditional Motion Compensation," filed Sep. 19, 2002, and U.S. Ser. No. 10/001,736, titled "Method and Apparatus for Adaptive Encoding Framed Data Sequences," each of which is incorporated herein by reference in its entirety.

[0056] To illustrate certain elements of the solution according to one embodiment, the energy-scalability of OFDM-based WLAN such as proposed by the 802.11a/g standards will be considered as an example. Currently proposed techniques to exploit the energy-performance trade-off in wireless links mainly rely on the reduction of the transmit-power. However, when considering the energy consumption breakdown of practical IEEE 802.11a WLAN transceivers (**FIG. 1**), it appears that the transmit-power itself accounts only for 8% of the total energy consumption of the transmitter. The Power Amplifier (PA) dominates with 49%. Moreover, due to the stringent linearity and spectral mask requirements of the OFDM modulation, class A power amplifiers are most often required. Consequently, the actual energy consumption of the PA is invariant with the transmit power. On the other hand, traditional receiver architectures are dimensioned for the most demanding modulation and present constant processing gain. Therefore, only the duty cycle effectively affects the link energy consumption. In that context, the highest modulation order that goes through the channel with reasonable packet error rate (i.e. 1%) is always the best choice. Performance (i.e. net data rate) downscaling does not bring effective energy consumption reduction.

[0057] In order to get effective energy consumption benefit from energy-aware link adaptation technique, the energy consumption scalability of OFDM-based WLAN transceivers has to be enhanced. The following enhancements are proposed: A. configurable power amplifier saturation point; and B. configurable receiver processing gain.

A. Configurable Power Amplifier Saturation Point

[0058] The transmitter energy-consumption is dominated by the power amplifier contribution. The idea is to allow adapting the gain compression characteristic of the power amplifier together with its working point on this characteristic. Typical WLAN power amplifiers operate with class A amplifiers that have a fixed gain compression characteristic, so that it is impossible to reduce simultaneously the output power and the linearity, e.g. to adapt to the lower requirement of lower order sub-carrier modulations, as illustrated in **FIG. 2B**. Reducing the output power translates into an increased back-off, and therefore into an higher linearity and a decreased power efficiency. This kills the benefit in terms of total power consumption. Therefore, it makes sense to vary independently the average transmit power and the back-off, which requires to adapt the gain compression characteristic (**FIG. 2A**). This can be done for instance by applying dynamic voltage scaling.

B. Configurable Receiver Processing Gain

[0059] For the receiver, it appears that the energy consumption is dominated by the digital signal processing in which the forward error correction accounts for an important part (**FIG. 1**). Scalability in receiving energy consumption can be achieved by trading off processing gain (and conse-

quently, receiver sensitivity) and computing energy by activating more or less digital signal processing. This can be achieved for instance by adopting iterative forward error correction (FEC) schemes like turbo-codes. Turbo-codes can be implemented with low energy consumption and provide a wide freedom to trade-off energy consumption versus code gain, which translates in receiving processing gain.

[0060] The impact of the energy-scalability enhancement techniques is now further analyzed. The optimal trade-off between the net data rate on top of the data-link layer (goodput) and the total link energy consumption is derived by exploring the settings of the traditional link parameters: modulation, code rate, output power, packet size next to the new functional parameters introduced, namely the power amplifier saturation power relative to the output power (back-off) and the number of decoding iterations. Table 1 summarizes the system level control knobs used to bring energy-scalability and their range.

TABLE 1

PA transmit Power (P_{Tx})	0 to 20 dBm by step of 2 dBm
PA back off (b)	6 to 16 dB by step of 1 dB
Packet (data) size (L_p)	50, 100, 500, 1000, 1500 bytes
Code rate (R_c)	½, ⅔, ¾
Modulation order (N_{mod})	1, 2, 4, 6
Number of turbo-iterations (I)	3 to 6

[0061] To be able to explore the energy-scalability of extended WLAN transceivers, robust performance and energy models are developed, which also consider the specificity of the indoor propagation channel. The models are incorporated in a simulation framework that allows generating very quickly performance-energy trade-off curves in realistic user scenarios, e.g. Pareto curves. These models are now discussed more in detail.

[0062] Tracking analytically the dependencies between the link parameters (system level knobs), the environment constraints (here the path loss and the channel frequency selectivity), the energy consumption and the performance (data rate) is difficult when the complete system is considered (i.e. not only the physical layer). Performance and energy consumption are non-linearly coupled e.g. by the automatic repeat request (ARQ) mechanism. To capture accurately those effects, system profiling using event-driven protocol simulation is most appropriate. Yet, the protocol simulator requires information about the radio-link, i.e. the packet error rate (PER), the gross data rate and the transmission energy per packet, which is also dependent on the system level knobs and the environment constraints.

[0063] At the physical layer level performance and energy models can be decoupled. An end-to-end simulation chain (**FIG. 3**) implementing the energy-scalability enhancement techniques as described before is used to identify the physical layer PER performance as a function of the system level knobs. First the data is encoded by a turbo-encoder. Then, the bit stream is modulated using OFDM according to the IEEE 802.11a standard. When transmitted, the signal is hampered by the non-linearity of the power amplifier. This effect is captured considering an additive distortion power, the distortion being assimilated to a noise source (D). The transmitter output signal-to-noise and distortion ratio (SINAD) is related to the back-off according to a law obtained by measurements. Next, the signal passes through the channel. The effect of the wireless indoor channel is decoupled in two effects: first, an average path loss assumed constant (a(d)); secondly, the multi-path fading that leads to a time-varying frequency selective channel response. To model the noise introduced by the receiver, a white Gaussian noise (n) is added. Finally, the signal is demodulated and decoded. The Packet error rate (PER) can be evaluated with Monte-Carlo simulations. However, it is better to assess the code block error rate (BIER), which is independent of the packet size. The PER can be directly derived from the BIER when knowing the packet size.

[0064] Simulation results depict a large variation in performance depending on the channel realization, even when the average path loss is constant. By performing Monte-Carlo simulations over a large amount of channel realizations, a classification has been made (**FIGS. 4A and 4B**). Channel responses leading to similar BIER for the same SINAD have been grouped in representative classes corresponding to so-called channel states. It turns out that 8 classes are sufficient to cover the range with representative average BIER vs. SINAD curves. **FIG. 4A** depicts the block error rate performance curves for the first 7 states considered. The histogram of the state probabilities is depicted in **FIG. 4B**. The separation between two adjacent channel classes is set to 2 dB, which corresponds to the granularity with which it is assumed the transmit power can be tuned. No transmission is assumed to be possible in the 8^{th} channel state, which occurs with a probability of less that 0.5%. With such a model, the energy-performance trade-off can be analyzed independently for each channel state, which again reduces the exploration complexity.

[0065] Next to the performance model, an energy model for the wireless link has to be established. The energy to transmit and to receive one packet is given as a function of the system level knobs (Table 2).

[0066] For the transmitter, the energy per packet is assumed equal to the product of the radio power by the time needed to transmit the packet (T_{up}), plus the energy consumed by the digital processing to format it, which is proportional to the packet size. E_{DSP} denotes the DSP energy consumed per transmitted bit. The radio power is the sum of the PA power (P_{PA}) and radio front-end power (P^{T}_{FE}). The PA power (P_{PA}) is equal to the transmit power (P_{TX}) divided by the PA power efficiency (η), which is expressed as a function of the back-off b by the measured relation $\eta(b)$. This relation has been obtained by fitting to measurement results. The contribution of the transmit front-end (P^{T}_{FE}) is assumed to be constant. The time to transmit a packet (T_{up}) can be computed as the ratio of the packet size (L_p) by the physical layer data rate, which is the product of the modulation order (N_{mod}), the number of data carriers in the OFDM symbol (N_c), the code rate (R_c) and the OFDM symbol rate, or baud rate (R_s).

TABLE 2

Performance model			Energy model		
			Model Parameters		
L_b	Turbo Code Block size (bit)	288	P^T_{FE}	Transmit front - end power (W)	200
W	Bandwidth (MHz)	20	P^R_{FE}	Receive front - end power (W)	200
T	Temperature (Celsius)	25	E^R_{DSP}	Receive DSP energy (nJ/bit)	7.5
NF	Receiver Noise Figure (dB) (Including Implementation loss)	15	E^T_{DSP}	Transmit DSP energy (nJ/bit)	8.8
			R_s	Baud Rate (kBaud)	250
d	Distance Transmitter - Receiver (m)	10	N_c	Number of data carrier per OFDM symbols	48
			Relations		
A(d)	Channel attenuation	From average propagation loss model	R_{bit}	OFDM bit rate (bps)	$R_{bit} = N_{mod} \cdot R_c \cdot R_s$
D(b)	PA distortion	Fitted on measurement on microsemi PA	T_{up}	Transmit/Receive on time per packet (s)	$T_{up} = \frac{L_p}{R_{int}}$
N	Receiver noise	$kT \cdot W \cdot NF$	$\eta(b)$	PA power efficiency	Fitted on measurement on microsemi PA
SINAD	Signal to noise and distortion ratio	$SINAD_c = \frac{P_{tc} \cdot A}{D(b) \cdot A + kT \cdot W \cdot NF}$	P_{PA}	PA power (W)	$P_{PA} = \frac{P_{pTz}}{\eta(b)}$
BIER	Code Block Error Rate	BIER = f(SINAD, Nmod. Re. 1) Fitted on Monte - Carlo Simulation Results	E_{FEC}	Energy turbo - decoding (J/bit)	$E_{FEC} = 3.98 \cdot 10^{-9} \cdot L_B + 1.21 \cdot 10^{-7} \cdot I + 1.32 \cdot 10^{-9} \cdot L_b + I - 3.69 \cdot 10^{-7}$
PER	Packet Error Rate	$PER = 1 - (1 - BIER)^{\frac{L_p}{L_b}}$	E_{TX}	Total Transmit energy (J/bit)	$E_{Tx} = (P_{PA} + P^T_{FE}) \cdot T_{up} + E^T_{DSP} \cdot L_p$
			E_{RX}	Total Receive Energy (J/bit)	$E_{Rx} = P^R_{FE} \cdot T_{up} + (E^R_{DSP} + E_{FEC}) \cdot L_p$

[0067] For the receiver, on the other hand, the energy per packet is modeled as equal to the analogue receiver power (P^R_{FE}) multiplied by the T_{up}, plus the digital receiver energy per packet, including the turbo-decoding energy expressed in function of the number of iterations and the code block size (N). The time to receive the packet is computed the same way as for the transmitter.

[0068] Table 2 summarizes the main relations and parameters of the performance and energy models. Parameter values have been captured from measurements carried out on a real advanced WLAN set-up.

[0069] As performance metrics the user data rate on top of the data link control layer (DLC) is considered and as energy consumption metric, the total energy to transmit successfully a packet. To be able to profile those metrics as a function of the system level knobs settings, considering the impact of the MAC and DLC protocols, the physical layer models are plugged into a network simulator. A two-user scenario has been defined. To evaluate the average user data rate (defined on top of the data link layer) and the corresponding energy consumption, the input queue of the transmitter is filled with 10000 packets and profile the total time and energy consumption needed to receive them correctly in the receiver output FIFO.

[0070] The simulation is carried out for every combination of the system level knobs described earlier. This leads to a variety of energy-performance trade-off points from which only the Pareto-optimal ones are retained. The latter points form a so-called optimal trade-off curve (also called Pareto curve) that gives the minimum energy to achieve a given performance in the considered channel state. Such trade-off curves are generated for standard and energy scalable systems on the 7 representative channel states. Results are depicted in **FIGS. 5A-5G**. The energy per transmitted bit is considered so that the results are independent of the link duty cycle. For the standard system, the 1-dB gain compression point is assumed to correspond to an output power of 36 dBm, thus a back-off of 16 dB for the maximum output power (20 dBm). Standard convolutional codes are considered according to the 802.11a standard. Modulation order, code rate and transmit power are assumed to be adaptive as in state-of-the-art link adaptation schemes. Packet size is also assumed to be adaptive, which reveals to have impact only on the worst-case channel states. For the energy scalable system, the additional system level knobs are varied according to the ranges given in Table 1. Both experiments have been carried out assuming an average path loss of 81 dB, corresponding typically to a distance of 10 m. From **FIG. 6**, it can be seen that with the standard system, it is always better to use the highest modulation order and code rate that go through the channel with reasonable packet error rate (i.e. 1%). This is due to the fact that the increase in energy per bit relative to the usage of a higher order modulation is compensated by the reduction in duty cycle consecutive to the increase in spectral efficiency. In that case, power management techniques based on duty cycle reduction are more effective.

[0071] The implication effect of 'lazy scheduling' on the end-to-end performance of a wireless network will now be analyzed. Further, it is investigated how the end-to-end performance versus energy can be effectively controlled by

adapting the link layer 'lazy scheduling' policy. The performance on top of the transport layer is considered. Indeed, it is the latter that actually provides the communication service to the application. Application specific performance metrics are not considered, but the TCP transport protocol is used, which is by far the most used in the Internet. The impact of variable error rate has already been analyzed (see L. Zou et al., 'The effects of Adaptive Modulation on the TCP Performance', Communications, Circuits and Systems and West Sino Expositions, pp. 262-266, 2002) but the consequence of the variable rate that would be introduced by 'lazy scheduling' has not. TCP throughput degradation resulting from varying rate and delay is discussed in Choi et al. ('TCP Performance Analysis in Wireless Transmission using AMC', IEEE VTC Spring, 2003) for CDMA mobile systems using adaptive modulation and coding. However, the possible control of this degradation and the trade-off with energy are not discussed in any of these related works.

[0072] From Shannon one knows that the minimum power P required to reliably transmit on a given (Gaussian) channel (characterized by a given signal to noise ratio and bandwidth) is a monotonously increasing, convex function of the targeted data rate R. This function is given by equation (1) where R_s is the symbol rate (baud rate), A and α the constant (average path loss) and variable (fading) components of the channel attenuation. N_o is the noise power density.

$$P = \frac{N_0 R_s}{A\alpha} \cdot \left(2^{\frac{R}{R_s}} - 1\right) \qquad \text{(equation 1)}$$

[0073] When the channel presents a time-varying attenuation, the signal to noise ratio varies accordingly and consequently the feasible rate for a given transmit power. A radio link control scheme is designed that allows finely controlling the link performance versus transceiver power consumption trade-off by adapting automatically, per frame, the discrete link control knobs introduced previously (modulation, code rate, transmit power and linearity) to tractable channel state information. Adapting the transmit rate and power to time-varying channel conditions in order to maximize the average data rate under an average power constraint is a well-understood problem referred to as 'link adaptation'. Optimal rate and power allocation schemes have been proposed. The dual problem, i.e. minimizing the average power under an average rate constraint can effectively be considered for power management. This can be seen as an application of the 'power aware' design paradigm where performance is traded off with energy. From the data link layer point of view (where normally rate and power control are implemented), the performance is traditionally evaluated in terms of net throughput (goodput), which is the net average data rate considering possible retransmissions. When a variable rate is considered, link delay becomes a second important performance metric. Indeed, as shown in **FIG. 6**, a queue has to be introduced to deal with the transmission rate variation. This queue introduces a delay that dominates the transmission latency at the bottleneck link. It can be shown (Little law) that the average queuing delay ($\bar{D}$) equals the average queue backlog ($\bar{A}$) divided by the average rate on the link ($\bar{R}$).

[0074] The energy versus queuing delay trade-off in such systems is extensively studied in '*Power and Delay Trade-offs in Fading Channels*', Berry, Ph.D. Thesis, MIT, Cambridge, Mass., December 2000. Long-term average power and queuing latency are considered. Using dynamic programming, policies to specify for each packet, depending on the actual queue backlog and channel gain, which rate to use are characterized. It is shown that queue stability can be guaranteed—i.e. the maximum number of bits in the queue is bounded—and the average energy versus average delay trade-off is bounded.

[0075] 'Lazy scheduling' is an example of such a policy. The principle of lazy scheduling consists of looking ahead at the packet arrivals, considering a fixed time window (T_w). At the end of each time window, the actual queue backlog (Δ_w) is considered to compute an average rate constraint (eq. 2), which is used to compute the transmit rate and power as a function of the channel attenuation for the next window. The 'water-filling in time' algorithm can be used for that purpose. Initially, this procedure has been derived considering the information theoretical relation between transmission rate and power (eq. 1). However, this relation corresponds to a bound that practical physical layer modulation and coding schemes tend to approach but do not meet. Also, in practice, the rate cannot be varied continuously but only stepwise, e.g. by varying the modulation order (constellation size) or the coding rate. Without hampering the generality, the practical rate adaptation algorithm proposed by Schurgers (cfr. supra) is considered. In an embodiment, the adaptation policy must indicate which modulation to use as a function of the channel attenuation. It is shown that for a narrow band Rayleigh fading channel, using Quadrature Amplitude Modulation (QAM) with constellation order $2^b=2^{2j}$, $j\in 1N$, the optimal policy is given by equation (eq. 3), where δ is obtained by solving (eq. 4), R_s being the symbol rate of the modulation and R corresponding to the average rate constraint. The corresponding power is given by (eq. 5) where C is a constant depending on the average path loss, coding gain, receiver noise figure and targeted bit error rate.

$$\bar{R} = \frac{\Delta_w}{T_w} \qquad \text{(equation 2)}$$

$$\begin{cases} b_i = 2 \cdot \left\lfloor \frac{1}{2} \cdot \log_2\left(\frac{\alpha_i}{\delta}\right) + 1 \right\rfloor & \text{if } \alpha_i \geq \delta \\ b_i = 0 & \text{if } \alpha_i < \delta \end{cases} \qquad \text{(equation 3)}$$

$$2 \cdot \sum_{j=0}^{\infty} \exp\left(-\delta \cdot 4^j\right) = \frac{\bar{R}}{R_s} \qquad \text{(equation 4)}$$

$$P_i = \frac{C}{\alpha_i}\left(2^{b_i} - 1\right) \qquad \text{(equation 5)}$$

[0076] Notice further that block fading is assumed, i.e. a constant channel attenuation during the transmission of one packet. This requires the channel coherence time to be high relative to the transmit duration of one packet. Hence, the modulation does not have to be adapted during a packet transmission.

TABLE 3

R_s	10 MBaud
C	60 mW

TABLE 3-continued

B	{1, 2, 4, 6}
P_{max}	20 dBm
Tc	20 ms

[0077] Considering this simple but representative model, the energy-delay trade-off achieved by lazy scheduling is evaluated. The channel state is varied according to a 8-state finite state Markov model whose transition matrix is set so that the probability density function (pdf) approximates a Rayleigh distribution of average 1 and the coherence time equals T_c. Further parameters are summarized in Table 3. The maximum average rate achievable with this setup, considering the transmit power limit of 20 dBm is close to 20 Mbps. In this experience, a constant input bit rate of 10 Mbps is considered, corresponding to 50% utilization of the link capacity.

[0078] Results are depicted in **FIG. 7** (circles). They are compared to the bound (optimal policy) computed for the same set-up (square). Clearly, the initial lazy scheduling performs sub-optimally when the tolerated latency is low (a short coherence time). To circumvent this drawback, instead of considering a finite time window, an adaptation of the average rate constraint with a finer grain is proposed. A proportional control is considered. The rate constraint is set proportional to the number of bits in the queue (eq. 6) at a frequency corresponding to the coherence time (T_c).

$$\bar{R}_i = K \times \Delta_i \quad \text{(equation 6)}$$

[0079] Here, $\bar{R}_i$ and Δ_i denote average rate constraint and average delay, respectively, i denoting a time index. The rate/power adaptation policy is tuned continuously according to this varying average rate constraint ($\bar{R}$) in (eq. 4). For a low proportionality factor (K), the system will tolerate a large average delay and react softly to the channel variation. Hence, the adaptation policy will behave close to the unconstrained water-filling in time. However, for high K, the reaction will be more aggressive, resulting in a lower average delay. For very high K; the system will always tend to select the highest modulation order possible and adapt the power according to channel inversion (i.e. boost the power for bad channels). Hence, by varying K, one can range between the two extreme link adaptation policies: water-filling in time and channel inversion.

[0080] It can be seen from **FIG. 7** (triangle) that the new policy beats the initial 'lazy scheduling' for smaller latencies and stays close to the hypothetic optimal policy in that region. Power-delay trade-offs for different input bit rates are depicted in **FIG. 8**. It can be seen that our adaptive policy offers a significant trade-off range for the average rate up to 80% link utilization. For half load (50% utilization), the average power ranges a factor 3 between the extreme adaptation policies. In **FIGS. 9A-91**, the average delay for each of the different input bit rate is plotted as a function of the parameter K. It clearly appears that this knob enables to control the average delay.

[0081] Interestingly, one can notice that the proposed policy is separable: the adaptations to the queue backlog—i.e. to the traffic requirements—and to the channel can be de-coupled as far as a "constraint propagation" is added. More specifically, an average rate constraint is derived from the queue backlog. This constraint is propagated to tune the rate/power versus channel state adaptation policy. Hence, the solution can be implemented without necessarily jointly designing the rate/power adaptation (in practice, the radio link control layer) and the data link control scheduler (that controls the queue). Those layers can be designed separately, the average rate constraint being the minimal cross-layer information to be passed from one layer to another to guarantee close-to-optimal operations.

[0082] So far, the energy versus performance trade-off resulting from the different 'lazy scheduling' approaches has been studied from the data link layer viewpoint only. Queuing delay and throughput on a single wireless link are considered as performance metrics. The different algorithms try to minimize the power needed to achieve these given performance metrics. Yet, to effectively study the impact of this technique on the end-to-end system performance, it is mandatory to also consider the interactions with the higher layers. Indeed, it should be understood how a bottleneck link (supposed to be the wireless link) delay and throughput translates into end-to-end system performance. The end-to-end system considered here consists of a single wireless link, corresponding e.g. to point to point scenario (**FIG. 10**). This is indeed the simplest scenario sufficient to show the cross-layer interactions. The single wireless terminal considered runs a number of applications, e.g. file transfers, which most probably use a TCP/IP protocol stack. The impact of an adaptive link layer on the average end-to-end TCP performance will now be shown. First the mechanism of the protocol is briefly reviewed. Secondly, it is shown how average end-to-end performance can be controlled by adapting the 'lazy scheduling' policy at link level.

[0083] TCP offers a reliable connection to the application. To enable this, acknowledgements are used to inform the source if a packet (identified with a sequence number) is well received. Using this feedback mechanism, it is possible to derive implicitly information about the possible network congestion, which occurs when the traffic sent through the network is larger than the capacity of the bottleneck link. Network congestion translates into excessive queuing delays or eventually packets drops at the bottleneck queue. Delays are discovered when the acknowledgement as a response to a sent packet is delayed more than expected by the source (i.e. a time-out event). Packet drops are translated in the reception of 'triple-duplicate' acknowledgements, i.e. the acknowledgements of packets following the lost packets contain identical sequence number. TCP reacts on this by maintaining a congestion window of W packets. Each Round Trip Time (RTT), i.e. the time between sending a packet and receiving its acknowledgement, TCP sends W packets. During congestion avoidance, the window is increased by 1/W each time an ACK is received. A TCP connection can also be in the slow start phase, where the window size is increased more aggressively. As one is mainly interested in the steady-state average behavior of TCP, this phase is not considered for the analysis. Conversely, the congestion window is decreased whenever a packets loss is detected, with the amount of the decrease depending on whether packet loss is detected by a duplicate ACK or by a time-out event. For a duplicate ACK, the window size is halved, while for a time-out it is reset to 1.

[0084] The steady-state performance of a bulk TCP flow (i.e. a flow with a large amount of data to send, such as file

transfers) may be characterized by the send rate, which is the amount of data pushed by the sender per time unit. If the number of packet losses or retransmission is small, the throughput, i.e. the amount of data received per time unit, is well approximated by this send rate. Define p to be the packet loss probability of a sent TCP packet. The send rate of a bulk TCP transfer is well approximated by:

$$S = \frac{1}{RTT}\sqrt{\frac{3}{2p}} \qquad \text{(equation 7)}$$

[0085] Considering a wireless link with time-varying rate, if the link adaptation is done well and MAC retransmissions are allowed, the losses at the wireless link can be neglected (to e.g. 0.1%). Hence, p is dominated by the losses or delays at the queue. Also RTT is mainly determined by the queuing delays at the bottleneck link (i.e. the wireless link). Therefore, both p and RTT depend largely on the link control parameter K. In **FIG. 11**, p denotes a range of K. A scenario with 10 TCP bulk transfer connections and a queue size of 128 packets is considered. The packet size is 1500 bytes. Further simulation parameters are the same as in Table 3. The loss probability p is approximated by determining a lower bound of the total number of packets lost or delayed. Hence, this gives a positive approximation of p. For small K, the throughput is determined by the link delay, which is mainly translated into time-out events. A larger K results in more aggressive TCP flows, filling the queue until a drop event, which is translated in more duplicate ACK events. The simulated TCP send rate per connection is plotted in **FIG. 13**, next to the value computed from equation (eq. 7). The shape of the send rate is indeed well approximated by (eq. 7), although it gives too optimistic values because of the optimistic approximation of p. Based on the link control parameter K, a useful range of throughput values is achieved, resulting in an end-to-end energy per byte versus throughput trade-off as depicted in **FIG. 12**. For readability, both throughput and energy axis are normalized to their maximum values (165 KB/s and 2 μJ/byte per connection). Energy scales in factor 2.5 when throughput scales in a factor 2, corresponding to K ranging from 2 to 10.

[0086] The above analysis shows that the end-to-end performance for bulk TCP transfer, i.e. the steady-state end-to-end throughput, is mainly determined by the queuing delay, translating into a loss probability p and the round trip time RTT. A performance constraint on top of TCP (throughput) can be translated (by eq. 7) into an average delay constraint on the wireless link, provided that the latter is the bottleneck. When lazy scheduling is considered, this delay can be controlled either by the look-ahead window size (original lazy scheduling proposal) or by the proportional control parameter (K) in the scheme according to one embodiment. Also, recall that the proposed link adaptation policy is separable. From the average delay constraints and the actual queue backlog, an average rate constraint can be derived (eq. 2 or 6) and propagated to the radio link control, which can use it to decide, for each packet, which rate and transmit power to use in order to minimize power. This observation has important consequences on the power management design paradigm. Indeed, this shows that efficient power management trading off energy with effective user-related performance metrics (here the TCP throughput) can be achieved keeping a clean, layered system architecture and its obvious advantage in terms of maintainability. Further, unintended interactions between protocols introduced by flat cross-layer design are avoided. Stack-wide power management is achieved by coordinating local algorithms by constraint propagation. The average DLC queuing delay and average transmission rate have been shown to correspond to the minimum information passing required between the transport layer and data link control, and between data link control and radio resource control, respectively. This allows to draft the structure of a stack-wide power management scheme distributed between the traditional layers (**FIG. 14**), still performing as well as a flat cross-layer power management scheme.

[0087] Referring to **FIG. 14**, the layers of the distributed stack-wide power management scheme, for example, include the following: application layer, middleware layer, transport layer (TCP), network layer (IP), data link layer (DL) and physical layer (PHY). The data link layer comprises radio link control and logical link control. With the distributed stack-wide power management scheme, one can conduct a method of managing the operation of a telecom system which comprises a queuing mechanism introducing a queuing delay. The method comprises: determining telecom environment conditions, which comprises current traffic requirements; determining an average rate constraint to meet the current traffic requirements, by solving an optimization problem and setting at least one control parameter, the control parameter taking into account the instantaneous queuing delay; and operating the telecom system by setting the control parameter taking into account the instantaneous queuing delay. In the method, the average rate constraint is set as a parameterizable function of the number of bits in the queue of the queuing mechanism. The parameters are control parameters. The determination of the average rate constraint is based on a look-ahead with a variable window size of link utilization. The window size is determined by said control parameter.

[0088] At run-time, the energy and performance model are calibrated (initially and whenever appropriate, still at a very low rate) to derive the actual energy/performance trade-off characteristics that are used to carry out the run-time settings adaptation at the desired (fast) rate. Parametric energy and performance models being available for the considered system, methods are now derived to carry out the run-time phase. The transmitter configuration is assumed to be done just before the transmission of the packet according to the channel conditions at that moment. The run time phase is split into two steps: a calibration (carried out at low frequency) and an effective run-time adaptation step (per packet). The latter relies on the calibrated energy/performance characteristics. Both steps are analyzed hereafter.

[0089] Knowing all parameters, included those measured at run-time (e.g. the average path loss), the energy per bit vs. goodput trade-off characteristic can be derived for each channel state. From the initial configuration space—i.e. the set of combination of control knobs settings (modulation order, code rate, transmit power and power back-off)—those corresponding to Pareto optimal trade-offs in the energy per bit versus goodput plane are stored in a table. This corresponds to the so-called calibration. At this point, the energy per bit is considered as energy metrics in order to maintain two independent axes. Indeed, for a given configuration, the

energy per bit is constant while the average power depends on the rate and on the duty cycle. If the knobs setting range is limited, an exhaustive search in the configuration space is still possible. If the number of knobs and/or the range is larger, then heuristics should be used. Notice that this search has to be performed only when the model is recalibrated (when entering the network or when the average path loss changes significantly), so the time and energy overhead involved in this is not really that critical. **FIG. 15** depicts the results of the calibration with the considered model for an average path loss of 80 dB, corresponding typically to a distance of 10 m indoor. One step-curve (Pareto curve) represents the Pareto-optimal energy-rate trade-off for a given channel state. Each of these points corresponds to a given (N_{mod}, R_c, P_{Tx}, b) setting. This implies that selecting one point on one of these curves corresponds to deciding jointly about the carrier modulation (N_{mod}), the code rate (R_c), the transmit power (P_{tx}) and the backoff setting (b) to be used.

[0090] The knobs settings corresponding to Pareto-optimal energy-rate trade-offs are known for each channel state from the calibration. Yet, this is not sufficient to carry out effective energy aware radio link control. A policy to decide which configuration point to use when a given channel state is detected, is needed. A trivial policy would select, on the Pareto curve corresponding to the current channel state, the point providing the goodput just larger than the user data rate requirement. Obviously, such a policy is sub-optimal from the power viewpoint. Indeed, since the constraint is the average goodput, it would be more effective to reduce the rate on the bad channel states (where the cost in energy to maintain a given rate is high) and compensate by increasing it on the good channel state (where the cost is lower). This is the principle underlying the water-filling algorithm proposed in '*The Capacity of Downlink Fading Channels with Variable Rate and Power*', A. J Goldsmith, IEEE Trans. Veh. Techn., Vol. 46, No. 3, August 1997. Yet, one cannot directly apply this algorithm here due to the discrete nature of the set-up. Therefore, first the structure of the problem is analyzed.

[0091] Let (r_{ij}, e_{ij}) be the coordinates of the i^{th} Pareto point on the curve corresponding to channel state j. The average power $\bar{P}$ and rate $\bar{R}$ corresponding to a given radio link control policy—i.e. the selection of one point on each Pareto curve—can be expressed by (eq. 8) and (eq. 9) where x_{ij} is 1 if the corresponding point is selected, 0 otherwise. ψ_j is the probability of the channel state j.

$$\bar{P} = \sum_j \psi_j \sum_i x_{ij} e_{ij} r_{ij} = \sum_i \sum_j x_{ij} \psi_j e_{ij} r_{ij} \equiv \sum_i \sum_j x_{ij} p'_{ij} \quad \text{(equation 7)}$$

$$\bar{R} = \sum_j \psi_j \sum_i x_{ij} r_{ij} = \sum_i \sum_j x_{ij} \psi_j r_{ij} \equiv \sum_i \sum_j x_{ij} r'_{ij} \quad \text{(equation 9)}$$

[0092] The notation p'_{ij} and r'_{ij} is introduced corresponding to the power and rate, respectively, when the channel state is j and the i^{th} point is selected on the corresponding curve, both weighted by the probability to be in that channel state. Only one Pareto point can be selected when being in one channel state, resulting in the following constraints:

$$\sum_i x_{ij} = 1, \forall\ j\ x_{ij} \in \{0, 1\} \quad \text{(equation 10)}$$

[0093] For a given rate constraint R, the optimal control policy is the solution of the following problem:

$$\min \sum_i \sum_j x_{ij} p'_{ij} \text{ subject to } \sum_i \sum_j x_{ij} r'_{ij} > R \quad \text{(equation 11)}$$

[0094] This is the classical multiple choice knapsack problem. One is interested in the family of control policies corresponding to R ranging from 0 to Rmax; Rmax being the maximum average rate achievable on the link. This family is called the radio link control strategy. Let kj denote the index of the point selected on the jth Pareto curve. Formally, kj=i ⇔xij=1. A control policy can be represented by the vector k={kj}. The control strategy is denoted {k(n)} corresponding to the set of point $\{(\bar{R}^{(n)}, \bar{P}^{(n)})\}$ in the average rate versus average power plane. A good approximation of the optimal radio link control strategy (i.e. that bounds the trade-off between $\bar{R}$ and $\bar{P}$) can be derived iteratively with a greedy heuristic explained in the flowchart shown in Table 4 below. Notice that this derivation is done when the energy-performance trade-off characteristics per state (the Pareto curves in **FIG. 15**) change, i.e. if recalibration occurs.

TABLE 4

Let $k^{(0)}$ be the solution of (11) when R = 0. Obviously it corresponds to $\{x_{0j} = 1\ \forall\ j\}$, i.e. the point of coordinate (0,0) is selected on each Pareto curve.

Denote the so-called current policy, $k^{(n)} = \{k_j^{(n)}\}$; the next policy $k^{(n+1)}$ corresponding to the just higher rate constraint can be computed as follows:

Compute the slopes $s_j^{(n)} = \dfrac{p'_{(k_j^{(n)}+1)j} - p'_{(k_j^{(n)})j}}{r'_{(k_j^{(n)}+1)j} - r'_{(k_j^{(n)})j}} \forall\ j$

Define $\hat{j}^{(n)} = \arg\max_j (s_j^{(n)})$

The next policy $k^{(n+1)}$ is such as $k_{\hat{j}}^{(n+1)} = k_{\hat{j}}^{(n)} + 1$ and $k_j^{(n+1)} = k_j^{(n)}, \forall\ j \neq \hat{j}$

[0095] From the family of policies derived with the greedy heuristic, it is possible to derive a policy for any average rate constraint ($\{\bar{R}^{(n)}\}$ also). For the given average rate constraint, from the set are selected those policies that lead to a average rate just higher (R_h) and just lower (R_l) than the constraint (R). For a given packet transmission, the configuration knobs are set according to one of these policies. The

selection is done randomly with probability P_h and $P_l=1-P_h$, respectively, where P_h is given by equation (12)

$$P_h = \frac{R - R_l}{R_h - R_l} \quad \text{(equation 12)}$$

[0096] Doing this, the final average power versus average rate characteristic is the linear interpolation between the trade-off points $\{(R^{(n)},P^{(n)})\}$.

[0097] The performance of the proposed radio link controller is evaluated in the average rate versus average power plane. Results are depicted in **FIG. 16**. The trade-off curve achieved by the control strategy derived with the greedy heuristic (circles) is compared to the optimal solution of (eq. 11) for the range of R (plain line). Theoretically, the greedy heuristic is sub-optimal since the Pareto curves in **FIG. 15** are not necessarily convex. Yet, it clearly appears that the optimal solution is closely approached. This trade-off curve can be used to evaluate the benefit of the proposed strategy compared with traditional radio link control aiming at maximizing the data rate under power constraint. With such a strategy, the policy corresponding to the highest feasible rate would be selected. This corresponds to the point A in **FIG. 16**. If the link capacity is not fully used, the link idles and the terminal enters regularly in sleep mode. The average power decreases linearly with the duty cycle. Hence, the corresponding power-rate trade-off is represented by the dashed-dot line in **FIG. 16**. It can be seen that, in all cases, except the one corresponding to maximal load, the proposed energy aware radio link control strategy saves power when compared with the traditional strategy. For example at half load, the gain reaches 80%.

[0098] Another embodiment relates to wireless communication systems that are provided with a sleep mode. Further also the multi-user aspect is introduced in the scheme to manage system-wide power management dimensions at runtime as described below.

[0099] In the wireless network as in **FIG. 17** multiple nodes are controlled centrally by an access point (AP). Each node (such as e.g. a handheld video camera) wishes to transmit frames at real-time and it is the AP's responsibility to assign channel access grants. In an embodiment, the resource allocation scheme within the AP must ensure that the nodes meet their performance constraints by delivering their data in a timely manner while consuming minimal energy. More formally the problem can be stated in the following way. The network consists of n flows $\{F_1, F_2, \ldots, F_n\}$. For notational simplicity, a one-to-one mapping of flows to nodes is assumed, but the design methodology is applicable to one or more flows per node. Each flow i, $1 \leqq i \leqq n$, is described by the following properties:

[0100] (a) Cost Function (C_i): This is the optimization objective, e.g. to minimize the total energy consumption of all users in terms of Joules/job. In for example a video context, a job is the frame size of current application layer video frame.

[0101] (b) QoS Constraint (Q_i): The optimization has to be carried out taking into account a minimum performance or QoS requirement in order to satisfy the user. As delivery of real-time traffic is of interest (e.g. video streaming), the QoS is described in terms of the job failure rate (JFR) or deadline miss rate. JFR is defined as the ratio of the number of frames not successfully delivered before their deadline to the total number of frames issued by the application. The QoS constraint is specified by the user as a target-JFR (JFR*), to be maintained over the lifetime of the flow.

[0102] (c) Shared Resource $\{R_1, R_2, \ldots, R_m\}$: Multiple resource dimensions could be used to schedule flows or tasks in the network, e.g. time, frequency or space. The restricted case is considered here where access to the channel is only divided in time. Therefore, time is the single shared resource and is denoted by R. The resource fraction consumed by the i^{th} node is denoted by R_i. The maximum time available for any flow is R_i^{max}, which is the frame period for periodic traffic.

[0103] (d) Control Dimensions $\{K_1, K_2, \ldots, K_l\}$: For a given wireless LAN architecture, there are platform independent control knobs or dimensions as already described before that control the received signal-to-noise ratio related to the resource utilization in terms of the transmission time per bit, given the current path loss. The control dimension settings are discrete, inter-dependent and together have a non-linear influence on the cost function.

[0104] (e) System state $\{S_1, S_2, \ldots, S_s\}$: As the environment is very dynamic, the system behavior will vary over time. The environmental factors independent of the user or system control are represented by a system state variable. In a wireless environment with e.g. VBR video traffic, the system state is determined by the current channel state and the application data requirements. The scheduling algorithm is executed periodically based on the channel epoch and the rate at which the data requirements change. Each flow is associated with a set of possible states, which determines the mapping of the control dimensions to the cost and resource.

[0105] The following system properties contribute to the structure of the methodology. The key aspects are the mapping of the control dimensions to cost and resource profiles respectively, and the general properties of this mapping. A resource (cost) profile describes a list of potential resource (cost) allocation schemes needed for each configuration point K. These profiles are then combined to give a Cost-Resource trade-off function, which may be essential for solving the resource allocation problem.

[0106] Cost Profile Properties

[0107] Every flow has a known minimum and maximum cost over all control dimensions, which is a function of the desired JFR* and the current system state (e.g. channel state). The cost range (difference between maximum and minimum) is important as this determines the impact the control dimensions have on the system cost.

[0108] The per-dimension discrete control settings can be ordered according to their minimal associated Cost.

[0109] The overall system cost, C, may be defined as the weighted sum of costs of all flows, where each flow can be assigned a certain weight depending on its relative importance or to improve fairness.

$$C = \sum_{i=1}^{n} w_i C_i \quad \text{(eq. 13)}$$

[0110] Resource Profile Properties

[0111] Every flow has a known minimum and maximum resource requirement across dimensions, which is a function of the desired JFR* and system state.

[0112] Depending on the current system constraints and possible configurations, each flow has a minimum resource requirement R_i^{min}. It is assumed all minimum resource requirements can be satisfied at each moment in time. Hence, no overload occurs and all flows can be scheduled. This is a reasonable assumption under average load and average channel conditions. However, under worst-case conditions and non-scalable video applications, a system overload may occur and one or more flows will need to be dropped. While the policy to drop flows is out of the scope of the optimization criterion, a practical system may employ a policing policy that is fair to the users.

[0113] The per-dimension discrete control settings can be ordered according to their minimal associated Resource requirement.

[0114] The overall system resource requirement R is defined as the sum of the per flow requirements:

$$R = \sum_{i=1}^{n} R_i \quad \text{(eq. 14)}$$

[0115] The goal is to assign transmission grants, resulting from an optimal setting of the control dimensions, to each node such that the per flow QoS constraints for multiple users are met with minimal energy consumption. For a given set of resources, control dimensions and QoS constraints, the scheduling objective is formally stated as:

$$\min_{C} \sum_{nodes} \omega_i C_i(S_i) \quad \text{(eq. 15)}$$

subject to:

$$JFR_i \le JFR_i^*,\ i = 1, \ldots, n \qquad (QoS \text{ Constraints})$$

$$\sum_{nodes} R_{i,j} \le R_j^{max},\ j = 1, \ldots, m \qquad \text{(Resource Constraints)}$$

$$K_j \rightarrow R_i(S_i),\ i = 1, \ldots, n;\ j = 1, \ldots, l \qquad \text{(Resource Profiles)}$$

$$K_j \rightarrow C_i(S_i),\ i = 1, \ldots, n;\ j = 1, \ldots, l \qquad \text{(Cost Profiles)}$$

[0116] The solution of the optimization problem yields a set of feasible operating points K, which fulfill the QoS target, respects the shared resource constraint and minimizes the system cost.

[0117] When considering energy-scalable systems, the number of control dimensions is large and leads to a combinatorial explosion of the possible system configurations as already explained. In addition, the resource and cost profile relations are complex. In order to solve this problem efficiently, a pragmatic scheme is needed to select the configurations at runtime. This is achieved by first determining the optimal configurations of all control dimensions at design time. At runtime, based on the channel condition and application load, the best operating point is selected from a significantly reduced set of possibilities.

[0118] A property of the design-time phase model is that the configurations can be ordered according to their minimal cost and resource consumption, describing a range of possible costs and resources for the system. For each additional unit of resource allocated, one only needs to consider the configuration that achieves the minimal cost for that unit of the resource. For each possible system state (for different channel and application loads), the optimal operating points are determined by pruning the Cost-Resource (C-R) curves to yield only the minimum cost configurations at each resource allocation point.

[0119] A function p_i: R→C is defined, such that

$$p_i(R_i(S_i)) = \min\{C_i(S_i) | (K_i \rightarrow R_i(S_i)) \hat{} (K_i \rightarrow C_i(S_i))\}$$

which defines a mapping between the Resource and the Cost of a certain configuration, k, for a node in a state, S_i, as shown in **FIG. 18**. Considering the resulting points in the C-R space, one is only interested in those that represent the optimal trade-off between the energy and resource needs for the system. Indeed, the trade-off between transmission time and transmission energy is a fundamental property for wireless communication, bounded by Shannon. Although the discrete settings and non-linear interactions in real systems lead to a deviation from this optimal trade-off, it can be well approximated as follows.

[0120] The convex minorant of these pruned curves is calculated along both for the Cost and Resource dimensions, and the intersection of the result is considered. As a result, the number of operating points is reduced significantly thus rendering it very useful for the runtime optimization (**FIG. 18**).

[0121] Several trade-offs are present in the system: increasing the modulation constellation size decreases the transmission time but results in a higher packet error rate (PER) for the same channel conditions and PA settings. In an embodiment, the energy savings due to decreased transmission time must offset the increased expected cost of retransmissions. Also, increasing the transmit power increases the signal distortion due to the PA. On the other hand, decreasing the transmission power also decreases the efficiency of the PA. Similarly, it is not straightforward when using a higher coding gain, if the decreased SNR requirement or increased transmission time dominates the energy consumption. Finally, considering the trade-off between sleeping and scaling: a longer transmission at a lower and more robust modulation rate needs to compensate for the opportunity cost of not sleeping earlier.

[0122] A greedy algorithm is employed to determine the per-flow resource usage R_i for each application to minimize the total system cost C. The algorithm traverses all flows' Cost-Resource curves and at every step consumes resources corresponding to the maximum negative slope across all flows. This ensures that for every additional unit of resources consumed, the additional cost saving is maximum

across all flows. The current channel state and application demand are assumed to be known for each node. This information is obtained by coupling the MAC protocol with the resource manager and is explained in the next section. By the assumption that sufficient resources are available for all flows, the optimal additional allocation to each flow, $R_i>0$, $1 \leqq i \leqq n$, subject to $\Sigma_{i=1}^{n} R_i \leqq R$ is determined.

[0123] The following greedy algorithm is used:

[0124] a. Allocate to each flow the smallest resource possible for the given state, R_{min}. All flows are schedulable under worst-case conditions, hence $\Sigma_{i=1}^{n} R_{min} \leqq R$.

[0125] b. Let the current normalized allocation of the resource to flow, F_i, be R_i, $1 \leqq i \leqq n$. Let the unallocated quantity of the available resource be R_l.

[0126] c. Identify the flow with the maximum negative slope, $|C_i'(R_i)|$—representing the maximum decrease in cost per resource unit. If there is more than one, pick one randomly. If the value of the minimum slope is 0, then stop. No further allocation will decrease the system cost further.

[0127] d. Increase R_i by the amount till the slope changes for the i^{th} flow. Decrement R_l by the additional allocated resource and increment the cost C by the consequent additional cost. Return to step b until all resources have been optimally allocated or when R_l is 0.

[0128] In this implementation, the configuration points at design-time are sorted in the decreasing order of the negative slope between two adjacent points. The complexity of the runtime algorithm is O(n.log L) for n nodes and L configuration points per curve. For a given channel and frame size, the number of configuration points to be considered at runtime is relatively small.

[0129] Taking into account that the relation $C_i(R_i)$ derived at design time is a convex trade-off curve, it can be shown easily that the greedy algorithm leads to the optimal solution for continuous resource allocation. The proof can be extended for real systems with discrete working points to show that the solution is within bounded deviation from the optimal. For a real system, however, the settings for different control dimensions such as modulation or transmit power are in discrete units. This results in a deviation from the optimal resource assignment. The worst-case deviation from the optimal strategy is bounded and small.

[0130] Again the practical example of an OFDM-based wireless LAN system is now considered. As previously discussed, several control dimensions can be identified that enable to trade-off performance for energy savings and vice versa. As above the power amplifier back-off ($P_{back\text{-}off}$), the power amplifier transmit power (P_{TX}), the modulation (N_{Mod}) and the code rate (B_c) are considered.

[0131] To determine the Job Failure Rate and total expected energy consumption, it may be essential to consider the system dynamics, which include the current channel condition and the application demand. The current job size (in number of fragments) varies significantly for video traffic. In addition, just considering the average received SINAD is not sufficient to characterize the channel for wireless OFDM systems where frequency-selective fading is important. A time-varying channel model is considered and expressions are derived relating the average packet error rate (PER), the JFR and expected energy consumption.

1) Traffic Model

[0132] As the goal is to provide timeliness (QoS) guarantees while minimizing energy consumption for a target performance, periodic delay-sensitive traffic is considered. Both constant bit rate (CBR) and variable bit rate (VBR) traffic is studied, in order to show the impact of the dynamics. VBR traffic consists of MPEG-4 flows. A Transform Expand Sample (TES) based MPEG-4 traffic generator that generates traffic with the same first and second order statistics as an original MPEG-4 trace is used. All fragmentation is done at the link layer and if a frame is not completely delivered to the receiver by its deadline, it is dropped. All applications employ UDP over IP.

[0133] Each frame size in fact maps to a different system state. A frame size is determined in a number of MAC layer fragments, which is assumed to be 1024 bytes long for this experiment. From the results, it is observed that for a given frame size, extrapolating the results for a curve within five fragments results in a very low approximation error. As the maximum frame size is assumed to be 50 fragments long in the tests considered, Cost-Resource curves are only constructed for 1, 2, 3, 4, 5, 10, 20, 30, 40, 50 fragments per frame.

2) Channel Model

[0134] A frequency selective and time varying channel model is now used to compute the PER for all transceiver knob settings. An indoor channel model based on HIPERLAN/2 was used for a terminal moving uniformly at speeds between 0 to 5.2 km/h (walking speed). This corresponds theoretically to a coherence time of approximately 28 ms. A set of 1000 time-varying frequency channel response realizations (sampled every 2 ms over one minute) were generated and normalized in power. Data was encoded using a turbo coder model and the bit stream was modulated using 802.11a OFDM specifications. For a given back-off and transmit power, the SINAD at the receiver antenna was computed as before. A path-loss of 80 dB at a distance of 10 m is assumed.

[0135] The signal was then equalized (zero-forcing scheme), demodulated and decoded. From the channel realization database, a one-to-one mapping of SINAD to receive block error rate was determined for each modulation and code rate. The channel was then classified into 5 classes (**FIG. 19A**). In order to derive a time-varying link-layer error model, we associate each channel class to a Markov state, each with a probability of occurrence based on the channel realizations database (**FIG. 19B**). Given this five-state error model, the PER can be efficiently modeled for different configurations at runtime. The Packet Error Rate (PER) is obtained by assuming the block errors follow a binomial process for a packet size of L_{frag} bits and a block size of 288 bits (see also Table 2):

$$PER=[1-(1-BlER)^{L_{frag}/288}] \quad \text{(eq. 16)}$$

[0136] Now the exact mapping of the control dimensions K to the cost and resource dimensions is derived, based on these expressions and the system state. When delivering a frame, different transmission and retransmission strategies can be used, each resulting in a different total expected energy and time to send the frame, and each resulting in another expected failure rate for the frame. To simplify, the policy is adopted that each fragment of a frame should be

transmitted or retransmitted using the same configuration K. This is a good approximation for the real optimal transmission strategy, which includes adapting the strategy depending on the outcome of a fragment transmission (conditional recursion which is complex to solve). For the approximation, a recursive formulation can be derived to compute the expected energy E_K, the timeslot needed $TXOP_K$, and the expected failure rate JFR_K, for each system state. The MAC protocol overhead is taken into account in this mapping, and the parameters, which are based on 802.11e, are listed in Table 2. A Contention Free access scheme is considered for the transmissions.

[0137] Consider a frame, which consists of m fragments or packets and has to be delivered during a known channel state C_S. The tuple (C_S,m) is defined as the system state. All following expressions for cost and resource consumption, hence, depend not only on the configuration K, but also on the system state. For notational simplicity, the state index is omitted. Each packet is transmitted with configuration K, for which the PERK can be determined, based on the models derived above. The probability that the frame is delivered successfully with exactly (m+n) transmissions (hence n retransmissions), is given by the recursion:

$$S_n^m(K) = \sum_{i=1}^{\min(m,n)} C_i^m \times (PER_K)^i \times (1 - PER_K)^{m-i} \times S_{n-i}^i \times S_{n-i}^i(K) \quad \text{(eq. 17)}$$

$$S_0^m(K) = (1 - PER_K)^m \quad \text{(eq. 18)}$$

in which C_i^m denotes the number of possibilities to select i objects out of m.

Hence, the probability to deliver the frame consisting of m fragments correctly with maximum n re-transmissions is

$$1 - JFR_n^m(K) = \sum_{j=0}^{n} S_j^m(K) \quad \text{(eq. 19)}$$

Here only data losses are assumed to result in job failures. As control frames are much shorter and less susceptible to errors, it is assumed they do not suffer from packet errors.

[0138] In order to determine the expected energy and time needed to deliver a frame with m fragments and n retransmissions, one needs to know the overhead due to the MAC protocol for a successful and a failed transmission, i.e. E_{good}, E_{bad}, T_{good} and T_{bad}. As a 802.11e HCF-type MAC is assumed, a successful and failed data transmission follow **FIG. 20**. Let E_{ACK} be the energy needed to receive an ACK packet, and T_{ACK} the time needed for the ACK, E_{Header} and T_{Header} are the overheads for the MAC and PHY headers.

$$E_{good}(K)=E_K+E_{Header}+(2\times T_{sifs}\times P_{Idle})+E_{ACK} \quad \text{(eq. 20)}$$

$$E_{bad}(K)=E_K+E_{Header}+((2\times T_{sifs}+T_{ACK})\times P_{Idle}) \quad \text{(eq. 21)}$$

$$T_{good}(K)=T_K+T_{Header}+(2\times T_{sifs})+T_{ACK} \quad \text{(eq. 22)}$$

$$T_{bad}(K)=T_{good}(K) \quad \text{(eq. 23)}$$

The time needed to send m fragments with max n retransmissions, for configuration K, is then:

$$TXOP_n^m(K)=[m\times T_{good}(K)]+[n\times T_{bad}(K)] \quad \text{(eq. 24)}$$

[0139] The average energy needed to transmit m fragments, with maximum n retransmissions, and configuration K is a bit more complex. The reason is that one is interested in the expected energy, taking into account the chance that a retransmission should happen or not:

$$E_n^m(K) = \sum_{j=0}^{n} S_n^m(K) \times ((m \times E_{good}(K)) + (j \times E_{bad}(K))) \quad \text{(eq. 25)}$$

[0140] It is the sum of the probability that the transmission will succeed after m good and j bad transmissions, times the energy needed for these good and bad transmissions. In order to have the correct expected energy consumption, a second term should be added to denote the energy consumption for a failed job, hence when there are less than m good transmissions, and (j+1) bad ones:

$$E_n^m(K) = E_n^m(K) + JFR_n^m(K) \times \left[E_{bad}(K) + \sum_{j=1}^{m} S_n^j(K) \times ((j \times E_{good}(K)) + (n \times E_{bad}(K)))\right] \quad \text{(eq. 26)}$$

[0141] As a result, the E, TXOP, and JFR can be determined as a function of frame size, channel state and number of retransmissions for each configuration K. This determines the full cost and resource profile for the system. In **FIGS. 21A and 21B**, the impact of the PA control knobs (PA back-off and PA transmit power) on the resource (TXOP) and cost (energy) is illustrated. Only the mapping that corresponds to the smallest TXOP and energy consumption is plotted (hence the mapping with the smallest number of retransmissions that achieves the required JFR* of 10e-3). **FIGS. 22A and 22B** shows the merged and pruned energy-TXOP curves for different channel states and different frame sizes, respectively. One can see that the total range in energy consumption is large, both within and across system states. The large trade-off illustrates the fact that traditional systems, which are designed for a fixed and worst cast scenario, result in a significant energy waste.

[0142] Determining a schedule that achieves the required performance with minimal energy consumption is challenging because the instantaneous load and channel state vary independently for each node. Previously the Energy and TXOP were determined needed to deliver a job depending on the current state, which is the frame size and the channel state. Based on these curves for each node, the scheduler can derive a near-optimal allocation at run-time. There needs to be a feedback loop between the nodes and the scheduler in the AP. The current state information needs to be collected by the scheduler and the decisions about the channel access grants should be communicated back to the nodes with minimal overhead. It is now shown how a sleep-aware Medium Access Controller can take care of this.

[0143] The MAC is responsible for resource allocation of the shared channel among different users. The packet-scheduling algorithm in the AP decides which node is to transmit, when, and for how long. The focus of frame scheduling is on the particular cases of frames sent from a

node to its associated AP (uplink) and also from one node directly to another and not via the commonly associated AP (peer-to-peer), as in **FIG. 17**. As all nodes are assumed to be within the transmission range of the AP, the AP can only schedule one peer-to-peer communication simultaneously. All nodes communicate only when they have a valid channel access grant or transmit opportunity (TXOP) or timeslot from the AP.

[**0144**] In order to instruct a node to sleep for a particular duration, the AP needs to know when the next packet must be scheduled. Waking a node earlier than the schedule instance will cause it to waste energy in the idle state. Waking the node later than the schedule instance, will cause it to miss the packet's deadline or waste system resources. The sleep-aware MAC protocol therefore employs two techniques to eliminate data dependency due to the application and channel. By buffering one frame, the AP can instruct the node to sleep from the time it was informed that the frame arrived to its scheduling instance. The AP still needs to poll the node upon frame arrival and therefore this only permits the node to sleep between the packet arrival instance and the packet schedule instance.

[**0145**] Buffering just two frames informs the AP of the current traffic demand but also the demand in the next scheduling instance. As shown in **FIG. 23**, the AP now only needs to communicate with the node at scheduling instances. As the real-time stream's packets are periodic, all idle time between transmission instances is eliminated for the ideal scheme.

[**0146**] In order to inform the AP of the instantaneous channel and the required application load for the current and next scheduling instances, one needs to employ a link-layer feedback mechanism. This is accomplished by adding just three bytes in the MAC header for the current channel state and the two buffered frame sizes. Protocols such as 802.11e provide support for channel information and queue sizes therefore require only minor modifications. In every transmission to the AP, the node communicates its channel state and packet sizes of the two heads of the line packets. In the acknowledgement, the AP instructs the node to sleep until the time of the next scheduling instance and also assigns it the duration of its next TXOP or resource slot. The scheduling decision is, thus, made every frame period (e.g. 30 ms for high-quality video) of the flow in the system with the highest frame rate.

[**0147**] The energy savings are now verified over a range of practical scenarios. For all results presented here, the target JFR* is set to 10^{-3} which is a reasonable value for wireless links. First the expected energy savings are analyzed across all channel states and the entire range of system loads. While this is profiled at design-time, it provides insight as to (a) the range and variation of energy savings with the system dynamics and (b) the contributions of energy savings from sleeping and scaling under different conditions. Following this, the energy savings at runtime are studied for both constant bit rate (CBR) and MPEG-4 video flows under different channel conditions and different system loads. In order to evaluate the relative performance of MEERA, the following comparative transmission strategies are considered:

[**0148**] 1. MEERA: This is the optimal operating scheme considering the energy trade-off between sleep and scaling. The operating point is determined from the C-R curves and the runtime algorithm described previously.

[**0149**] 2. MEERA-no sleep: This scheme uses the C-R curves to determine the optimal TXOP when no sleeping is possible. The same runtime algorithm is used and the nodes remain in the idle state after completion. The purpose of this case is to show the contribution of sleeping under different loads.

[**0150**] 3. Fixed: In this scheme the back-off and output power are fixed to the highest setting and use take the highest feasible modulation and code rate that will successfully deliver the packets. After successful transmission, one goes to sleep. This case is the approach where no scaling is employed and the sleep duration is maximized.

[**0151**] 4. Fixed—no sleep: This scheme is similar to Fixed, but the transceiver remains in the idle mode after successful transmission. This is the base operating scheme of current wireless LAN transceivers with no power save features enabled.

[**0152**] Consider the C-R curves in **FIG. 22**, for transmitting a one-fragment sized frame over different channel states. If the TXOP assignment for the user is distributed uniformly over time, the expected energy consumption is proportional to the area under the $C_i(R_i)$ curve for each channel state shown in **FIG. 22**(*a*). Similarly, the C-R curves are derived for the three comparative transmission strategies discussed above. The $C_i(R_i)$ curve, for the fixed case, for example, is in fact a fixed point (at small TXOP), and the expected energy consumption proportional to the area under a horizontal line through this point. In **FIG. 24A**, the relative energy consumption (normalized by the maximum energy consumed by Fixed over all cases), is plotted for the four schemes over different channel states. As expected, MEERA outperforms the other techniques since it takes advantage of the energy that can be saved by both sleeping and TXOP scaling. The energy needed to transmit a unit of data increases from best to worst channel state due to a combination of (a) the lower modulation rate necessary to meet the higher SINAD requirement (hence shorter sleep duration), (b) a higher required output power to account for the worse channel and (c) the increased cost of retransmissions. It can be observed, for example, for the best channel state, the energy consumption is low for both the Fixed and MEERA approaches. The energy gains for this channel state primarily result from sleeping. On the other hand, for the worst channel state, the transmission energy becomes more dominant, the energy gains due to TXOP scaling are more significant. When looking at the energy gains contributed by sleeping and scaling over a range of link utilizations, the energy gains due to TXOP scaling can be expected to dominate when the transmission duration is large when compared to the sleep interval. Hence, for larger frame sizes or at higher per-flow link utilization, the relative energy saving due to scaling has a grater influence. This observation is illustrated in FIG. **24**B, where the relative gain for the different techniques—compared to the Fixed-no sleep case—are plotted over a series of frame sizes (in terms of number of 1024 byte fragments), for channel state 1. Indeed, for a 5-fragment frame TXOP, scaling and sleeping (represented by MEERA-No sleep and Fixed-sleeping respec-

tively) have a comparable performance. For larger frame sizes, the TXOP scaling in MEERA-No sleep contributes significantly to the energy savings.

[0153] Now a multiple-user scenario is considered where the TXOP assignments are not uniformly distributed but based on the user's application data rate requirement and the constraints enforced by other users sharing the same link. The influence is now discussed of the aggregate link utilization on the per-flow energy consumption for CBR flows over a static channel. The effective throughput of 802.11e, after considering protocol overheads for the first channel state, is approximately 30 Mbps when the highest modulation constellation is used. In the experiment illustrated by **FIG. 25A**, the link utilization is increased in steps of 2 Mbps CBR flows up to the maximum link capacity. The per-flow energy consumption of MEERA increases as the aggregate system load increases. At higher loads due to a large number of flows, a smaller TXOP from the C-R curve is assigned to each flow resulting in higher per flow energy consumption. The difference with the case where only TXOP scaling is used (MEERA-no sleep) is most noticeable since the possibility to scale is reduced with increasing system load. In a multi-user scenario it is always beneficial to include sleeping as it is influenced to a lesser extent by aggregate load increases due to other flows. On the other hand, if the load is increased by increasing the per-flow throughput requirement for a fixed number of flows, the gain of scaling dominates over sleeping. In **FIG. 25B**, the average per-flow data rate requirement is increased for an MPEG-4 flow (in a network with five such flows). In this case, the reduction of sleep duration increases the energy consumption only slightly as the energy saving from scaling start contributing more. This is evident when comparing MEERA with the energy consumption of Fixed, which increases almost linearly with the data rate. It is important to note that as the system is not forced into overload, only moderate link utilization is considered (<70% average link utilization using the highest transmission rate) for MPEG-4 flows.

[0154] Now the energy consumption trends are considered for a time-varying channel. A 5-user scenario is used to understand the impact of dynamic channel variations on energy consumption. The channel varies independently over all the users on a frame-by-frame basis. In **FIG. 26A**, the total system load is increased from 2.5 Mbps to 10 Mbps for five CBR flows. When compared to the same system load where the channel is in the best (fixed) state, an increase in energy consumption is observed. This is because during every scheduling cycle, the flows experiencing worse channel states require more transmission time (due to lower constellation) and therefore consume more energy. In addition, they force the other flows to transmit in a smaller TXOP and increase their energy consumption too. Further, in **FIG. 26B** the energy savings achieved when compared to Fixed-no sleep over different system load are shown. Sleeping results in an energy saving for all loads, which is obvious in a multi-user scenario. However, comparing the gain MEERA achieves to the Fixed approach (with sleeping), MEERA results in a significant additional gain of a factor 2 to 5 depending on the system load. The combination of sleep and scaling in MEERA yields an overall system gain factor from 2 to 9.

[0155] **FIG. 27** is a flowchart of a process for managing the operation of a telecom system in accordance with another embodiment. According to the embodiment, the process comprises determining a rate constraint and a telecom environment condition (S**2701**) and providing predetermined working points per discrete set of telecom environment condition (S**2703**). The process may further include determining the predetermined working points. Then, a working point is selected from the predetermined working points (S**2705**), and the telecom system is operated at the selected working point (S**2707**). Prior to the selection of the working point, the process may further include loading and adapting the predetermined working points.

[0156] The optimization process can lead to much better solutions if all the factors are taken into account in a cross-layer approach. This way, we explore the cross-layer interdependencies and joint optimizations for varying network and video coding settings.

[0157] In one embodiment the application of this approach is demonstrated for a single user, single-hop wireless link, where the transmission occurs from the AP (Access Point) to the mobile device. But the generic concept is applicable directly in multi-user TDMA based contexts as well. In an exemplary embodiment we have also based the entire demonstration on industry-accepted standards for the layers, i.e. MPEG4 main profile, UDP, IP (Internet), WLAN 11a/e with a DCF or HCF scheduler. Other standards may also be used.

[0158] A main contribution is the identification of important latency-energy trade-offs that can be exploited for any multi- or single-user service involving interactive passing of data. In that case the user(s) can exchange a longer duration of their battery with a larger response time for the interaction over the wireless network (including the application layer). These trade-offs are only exposed through a full cross-layer modeling and analysis of the entire protocol stack. We have also introduced several run-time controllable knobs to enable this trade-off.

[0159] Some embodiments deal with battery-powered multimedia (in particular a video encoding system) combined with a telecommunication system (or services) (in particular a wireless system, with either single or multiple users, with direct (single-hop) or indirect (multiple-hop) communication. Some embodiments provide an energy-latency management system (and even energy-latency—Quality management system) for such combined system, said management system including storage of predetermined latency-energy trade-off's, where from said management system selects on the fly (run-time) a working point.

[0160] In one embodiment the settings for the video codec are fixed for the whole sequence.

[0161] In another embodiment an on-line video adaptation to the network condition is used.

[0162] Some embodiments have multiple predetermined operating points, classified to external conditions (for instance the channel state), said predetermination being done by performing optimization across the various layers, meaning by exploring at least a parameter from the multimedia application and a parameter from the wireless transmission system, more especially dealing with parameters of several layers of the wireless transmission system, like the transport, network and MAC layer.

Test Bench Definition

[0163] Transmission of the video data through the wireless channels and error corruption is simulated by means of a significantly extended version of the Network Simulator (NS-version 2). In the IMEC extension of NS the protocols, which are emulated are the following: UDP and IP version 6 for the transport, network protocols and at the MAC layer, the 802.11a and the 802.11e protocols (see International Standard for Information Technology—Telecom. and Information exchange between systems—Local and Metropolitan Area Networks—Specific Requirements—Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications, IEEE 802.11 WG, ISO/IEC 8802-11: 1999(E) IEEE Std. 802.11, 1999, and IEEE 802.11 WG, IEEE 802.11e/D8.0, "Draft Amendment to Standard for Telecommunications and Information exchange between systems—Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Quality of Service (QoS) Enhancements", February 2004.) where the Distributed Coordination Function (DCF) and the Hybrid Coordination Function (HCF) are simulated. For the HCF several scheduling periods are used: 8, 17 and 34 ms to cover a wide latency range, useful for interactive services like Vo IP or video games. The distance between AP and mobile device is taken as 10 m. But obviously, similar results can be obtained for other common distances in WLAN contexts.

[0164] The video codec used is a proprietary MPEG-4 codec at Simple Profile (I, P frames, no B frames) optimized towards memory, (for further detail see "Motion estimation for low power video devices", Proceedings of the IEEE International Conference on Image Processing (ICIP01), Thessaloniki, Greece, October 2001, pp. 953-957. C. De Vleeschouwer and T. Nilsson, and "Cost-efficient C-level design of an MPEG-4 video decoder", Proceedings of the IEEE Workshop on Power and Timing Modeling, Optimization and Simulation, pp. 233-242, Goettingen, Germany, September 2000. K. Denolf, et al) where error resilience is achieved by using resynchronization video packets of fixed size and random Intra Macroblock refreshes. In the NS each video packet is encapsulated in a separate UDP packet, which in turn is encapsulated in a single MAC packet. This way, the loss of a single MAC packet causes the loss of one video packet, being all video packets independently decodable. Of course also other packeting can be investigated.

[0165] The knobs used at the application layer are the quantization parameter (QP), which determines the encoding accuracy, the synchronization video packet size and the number of forced Intra Macroblocks per frame. A lower QP value corresponds with a higher encoding accuracy and therefore both higher bit rate and quality. The use of resynchronization video packets, also called slices, allows the decoding process to resynchronize and continue the decoding process when transmission errors occur. To overcome the degradation caused by these transmission errors, the introduction of Intra Macroblocks can be enforced in order to refresh information, which cannot be predicted from previous erroneous frames.

[0166] Nevertheless, in our study we observed that the percentage of Intra Macroblocks was not a useful knob as the targeted PER are so low that the refreshment of frames by means of a small percentage of Intra Macroblocks becomes sufficient and the results are shown for 8% of Intra MB refreshment.

[0167] It is also important to note that in the current demonstration the settings for the video codec are fixed for the whole sequence, without performing at this stage an on-line video adaptation to the network conditions.

[0168] Of course a run-time control loop to steer these knob settings based on the channel and user data scenario's is possible also.

[0169] At the network side the main knobs we can play with are the following: modulation type, code rate, transmission power, back-off power, and Medium Access Layer (MAC) packet size among others.

[0170] In our case the selection of the video packet size coincides with the MAC packet size as each video packet is encapsulated in a separate MAC packet, so that a MAC packet loss does not affect the decoding of other independently decodable video packets, which go in other MAC packets.

[0171] There is a tradeoff in the selection of the packet size both at the application and link layers. At the application side smaller resynchronization packets allow a higher chance of resynchronization in case of errors but also introduce a higher overhead and reduce the prediction efficiency within a frame. At the MAC layer smaller packets increase the overhead (and with it the transmission energy) but they also reduce the possibility of a packet being hit by a transmission error, which relaxes the protection requirements.

[0172] To fully explore the cross layer interdependencies we run tests not only with different video coding settings (varying QP, video packet size and percentage of intra macroblock refreshments) but with several configuration of the network parameters (code rate, constellation, packet size, transmission power . . .) Table 5 shows the range of the multimedia settings, which is taken into consideration. The range of network settings can be found in EP 1467294. The test bench used includes the Foreman sequence (medium degree of motion and bit rate requirements) and the Mobile sequence, which has a higher degree of motion and it is more demanding in terms of bit rate.

[0173] In addition the tests are performed for 8 different channel states to evaluate the impact of different transmission conditions and allow adaptability to those. These states have been shown to characterize the channel behavior well.

TABLE 5

Application knobs	
Sequence	Foreman, Mobile
QP	6, 7, 8 and 9
Packet size (bits)	4000, 6000, 8000, 12000 and 16000
Intra MB refreshment (%)	3, 8, 20 and 25

[0174] The NS-based simulator provides as a result for each network and video settings the energy consumed at both network source and destination, together with the maximum and average latency per video packet that the

network introduces. The latency introduced at the application level together with the energy consumed at the encoding/decoding process is added to the network values.

[0175] The energy consumed at the application for the encoding and decoding process of bitstreams with varying settings is measured with a tool. The processor energy is estimated to be comparable to the energy spent on memory accesses. This estimate gives only relative accuracy (not in absolute terms) but for our main conclusions, this relative accuracy is sufficient to substantiate the concept. The maximum encoding/decoding latency per packet is measured on a Pentium IV at 2.8 Ghz, where the encoding/decoding processes are downscaled to real-time processing (30 frames per second in the case of Mobile and 25 frames per second for Foreman). For other processor platforms, the absolute values will again change, but the relative differences are expected to hold.

[0176] Video quality is measured with the Structural Similarity Index (SSIM) (see "Video Quality Assessment Based on Structural Distortion Measurement" Signal Processing: Image Communication, vol 19, No. 1 January 2004, Zhou Wang, Ligang Lu and Alan C. Bovik). We have observed that as PSNR, the SSIM value also does not correlate properly with the subjective video quality, when transmission errors occur. This is the reason why the quality axis is currently rather taken as a constraint and high video quality is imposed for all Pareto points, corresponding then to Packet Error Rates equal to zero or very low so as to keep the perceived quality high.

[0177] In an alternative setting this constraint is relaxed and latency-energy-quality trade-offs are generated.

Simulation Results

[0178] Some benefits of the HCF scheme can be found in Appendix A as "Optimal Fixed and Scalable Energy Management for Wireless Networks".

[0179] FIGS. **28** to **43** show certain tradeoff points obtained in the Energy-Latency axes for channel states 2 and 4, for transmission modes DCF and HCF. These tradeoff points, so called Pareto points, are extracted for each of the video coding settings (combination of QP and packet size). As mentioned previously, the quality axis is not shown as good video quality is taken as a constraint. This way, all resulting Pareto points correspond to high video quality where the Packet Error Rate (PER) is zero or very low.

[0180] The general trends, as expected, are that the higher the bit rate of the encoded sequence (corresponding to lower QPs with a higher encoding accuracy) the higher the energy invested at both network (higher amount of packets to be transmitted) and application sides. At the same time, for a constant bit rate a higher packet size implies lower energy and latency at the network side, as the overhead is lower. On the other side, when the transmission conditions get worse, big packet sizes have a higher probability of being hit by an error, introducing then more errors at the link level, which decreases in turn the energy efficiency. In our case, as the transmission distances are relatively small and the network load is low (one user) so the channel conditions are not bad enough to show this effect.

[0181] As the channel conditions get worse the more energy needs to be invested at the network side to successfully transmit the video packets with a low or zero PERs. This can be observed when comparing channel states 2 (relatively good channel) and 4 (worse channel conditions). This way, for different channel conditions the optimal points in terms of network and codec settings vary.

[0182] At the application side, lower QP values, corresponding to higher encoding accuracy) increase the encoding/decoding energy. But in terms of encoding/decoding latency, for a fix packet size, the encoding/decoding latency per packet is higher for higher QPs. This is due to the fact that the encoding/decoding process is scaled down to real-time processing and then the higher the bit rate (lower QP) the higher number of packets per frame are used and their relative encoding latency decreases.

[0183] All graphs show that there is always a large tradeoff present between energy and latency. To achieve low packet latencies a higher energy needs to be invested at the network, while if we target lower energies the global latency will increase. By playing with both multimedia parameters (QP, packet size . . .) and network parameters (code rate, modulation, transmission power . . .) the amount of possible tradeoffs increases, enlarging then the Energy and Latency ranges. These enlarged ranges increase in turn the flexibility of the global system, which has more possible operation points and allows higher savings in energy or latency by playing with the proper knobs. The application of DCF and HCF modes of operation creates several Pareto points for each particular video setting. This, together with the use of several application knobs (video settings) allows to extract a global Pareto curve with extended energy and latency ranges.

[0184] In addition, the channel state highly influences the resulting tradeoffs as well as the network and application settings of all Pareto points in the global Pareto curve.

[0185] The energy-latency ranges that are obtained are often of factor 2 or 3, reaching even factor 4 for worse channel conditions (channel 4) and Mobile sequence while factors 1.3 to 2 are mainly obtained for Foreman.

[0186] It can be observed how the tradeoffs are larger for the Mobile sequence than for the Foreman, this is due to the larger bit rate variation among the extreme settings of Mobile (from 2 to 4 Mbps) than the one in Foreman (from 0.7 to 1.1 Mbps). Therefore, larger bit rate ranges imply larger energy and latency ranges. Similarly, the range of the energy-latency tradeoffs also increases with worse channels. Both effects of increasing the bit rate or having worse channel conditions are similar in the sense of demanding higher requirements from the system.

[0187] These large trade-off ranges can be exploited by a run-time controller that sets the appropriate knob positions based on the user QoE requirements. If battery power is low, latency can be sacrificed for the interactive services. In the other case, the user will probably prefer lower latencies. Without these cross-layer analysis and optimizations, the trade-off ranges would not be present. Note that even the "reference" points in the graphs below already contain the cross-layer interaction between the multi-user scheduler and the WLAN baseband and front-end layers, where a run-time controller is already steering the appropriate knob positions based on channel conditions. If the latter cross-layer optimization is removed, the remaining trade-offs will become minimal.

[0188] In the following graphs (**FIG. 28** to **43**) different colors correspond to different QP, while the point types correspond to different video packet sizes as follows:

[0189] QP 9: green

[0190] QP 8: cyan

[0191] QP 7: magenta

[0192] QP 6: blue

[0193] Packet size 4000 bits: '+'

[0194] Packet size 6000 bits: 'x'

[0195] Packet size 8000 bits: 'square'

[0196] Packet size 6000 bits: 'diamond'

[0197] Packet size 6000 bits: 'circle'

[0198] In **FIGS. 40 and 42** the green curve represents the Pareto curve for the bitstream configuration selected as reference (QP 9, packet size 8000 bits) while the pink curve represents the global Pareto curve with the best tradeoffs Energy-Latency among bitstreams of different settings in terms of QP and packet size.

[0199] Cross-layer adaptation including the multi-media application layer proves to be beneficial as we can observe that for the same latency value it is possible to obtain savings of 25% in the energy, while the quality is maintained or even increased. In the same way, lower latencies can be achieved by changing the video parameters (QP, packet size). These optimal operation points are dependent on the channel state, which is being considered. Note that all these points already include the cross-layer optimization (and run-time controller) between the multi-user scheduler and the WLAN baseband/front-end (as discussed above).

[0200] Tables 6 to 9 show the energy and latency values for points in both curves corresponding to low, medium and high latencies within the given range and for channel states 2 and 4.

TABLE 6

Reference Curve for Channel State 2.

Reference Curve	Wireless Setting	Latency (ms)	Energy (Joules)	QP	Packet size (bits)
High Latency	98	42.12	6.87	9	8000
Medium Laten	109	37.28	7.6	9	8000
Low Latency	NA	NA	NA	NA	NA

[0201]

TABLE 7

Optimal Curve for Channel State 2.

Optimal Curve	Wireless Setting	Latency (ms)	Energy (Joules)	QP	Packet size (bits)
High Latency	98	42.12	6.87	9	8000
Medium Laten	108	35.44	7.19	9	6000
Low Latency	106	27.5	10.02	6	8000

[0202]

TABLE 8

Reference Curve for Channel State 4.

Reference Curve	Wireless Setting	Latency (ms)	Energy (Joules)	QP	Packet size (bits)
High Latency	97	50.77	12.10	9	8000
Medium Laten	120	37.28	21.25	9	8000
Low Latency	NA	NA	NA	NA	NA

[0203]

TABLE 9

Optimal Curve for Channel State 4.

Optimal Curve	Wireless Setting	Latency (ms)	Energy (Joules)	QP	Packet size (bits)
High Latency	49	50.81	9.16	8	6000
Medium Laten	107	37.27	15.85	8	6000
Low Latency	125	24.69	27.1	6	6000

[0204] While the above detailed description has shown, described, and pointed out novel features of the invention as applied to various embodiments, it will be understood that various omissions, substitutions, and changes in the form and details of the device or process illustrated may be made by those skilled in the art without departing from the spirit of the invention. As one example, it will be appreciated that the test voltage driver circuitry could be separate from the array driver circuitry used to create the display. As with current sensors, separate voltage sensors could be dedicated to separate row electrodes. The scope of the invention is indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed is:

1. A method of managing the operation of a system comprising a processing subsystem configured to run a multimedia application and a telecommunication subsystem, the method comprising:

determining telecom environment conditions;

selecting a working point from a plurality of predetermined working points, wherein the selecting is based at least in part on the determined environmental conditions, the working points having been determined by simultaneously optimizing control parameters of both the multimedia application and the telecommunication subsystem;

setting control parameters in the multimedia application and/or the telecommunication subsystem to configure the system to operate at the selected working point; and

operating the system at the selected working point.

2. The method of claim 1, wherein part of said working points differ in that they correspond to different control parameters to be set on both said multimedia application and said telecommunication subsystem.

3. The method of claim 1, wherein said telecommunication subsystem comprises a single telecom device.

4. The method of claim 1, wherein said telecommunication subsystem comprises a plurality of telecom devices.

5. The method of claim 4, wherein operating said telecommunication subsystem comprises setting control parameters corresponding to said selected working point for each telecom device.

6. The method of claim 4, wherein said telecom environment conditions comprise at least one of path loss and channel frequency selectivity.

7. The method of claim 4, wherein the selecting of a working point comprises solving a scheduling problem.

8. The method of claim 7, wherein the scheduling problem involves scheduling transmission of packets between two or more of the telecom devices, wherein the scheduling takes into account dependencies between packets.

9. The method of claim 1, wherein the telecom environment conditions comprise one or more channel states.

10. The method of claim 9, wherein performance-energy trade-off curves are derived for each of the one or more channel states.

11. The method of claim 10, wherein the performance-energy trade-off curves comprise Pareto-optimal curves.

12. The method of claim 10, wherein energy-per-bit is used as an energy metric in the performance-energy trade-off curves.

13. The method of claim 10, wherein latency is used as a performance metric in the performance-energy trade-off curves.

14. The method of claim 10, wherein a sum of energy consumption of the telecom devices is used as an energy metric in the performance-energy trade-off curves.

15. The method of claim 1, wherein the plurality of predetermined working points define a monotonic, non-convex curve.

16. The method of claim 1, further comprising determining the plurality of predetermined working points.

17. The method of claim 16, wherein determining the plurality of predetermined working points is based on the determined telecom environment conditions.

18. The method of claim 1, further comprising loading the plurality of predetermined working points.

19. The method of claim 18, further comprising adapting the plurality of predetermined working points.

20. The method of claim 1, wherein the plurality of predetermined working points comprises a sleep mode and a working mode of said telecom device.

21. The method of claim 1, wherein said telecommunication subsystem is a wireless telecommunication subsystem.

22. The method of claim 1, wherein the control parameters to be set comprise at least one of controlling modulation order, code rate, transmit power, packet size, quantization parameters, and synchronization video packet size.

23. The method of claim 1, wherein the plurality of predetermined working points corresponds to a discrete set of telecom environment conditions.

24. The method of claim 23, wherein the discrete set of telecom environment conditions is organized per channel state.

25. The method of claim 1, wherein the selecting of a working point comprises solving an optimization problem.

26. A device operable according to the method of claim 1, the device comprising a memory configured to store performance/energy trade-off curves.

27. A device operable according to the method of claim 1 wherein the selecting of a working point comprises solving an optimization problem, the device comprising a processor configured to solve the optimization problem, or to determine or adapt the predetermined working points.

28. A computer readable medium storing a program configured to execute a method of managing the operation of a system comprising a processing subsystem configured to run a multimedia application and a telecommunication subsystem, the method comprising:

determining telecom environment conditions;

selecting a working point from a plurality of predetermined working points, wherein the selecting is based at least in part on the determined environmental conditions;

setting control parameters in the multimedia application and/or the telecommunication subsystem, wherein the setting of control parameters is based at least in part on the selected working point; and

operating the system at the selected working point.

* * * * *