IEEE 802.11
Wireless Access Method and Physical Layer Specifications

Title: The potential of Dynamic Power Control

Presented by:

Wim Diepstraten
NCR SE-Utrecht
NCR/AT&T Network Product Group
Nieuwegein
The Netherlands
31-3402-76482 (V)
31-3402-39125 (Fax)
Wim.Diepstraten@Utrecht.ncr.com (Email)

Issue's addressed:		
	Medium re-use efficiency strategy.	(??-?)
	What is the strategy for capacity control	(19-7)
	Must the MAC work on a single channel PHY	(25-3)
	Requirements for the MAC/PHY interface.	(12-2)
	Is MAC/PHY exchange needed for Management	(13-5)
	What is the MAC Frame structure	(20-4)

Abstract: This paper describes a Dynamic Power Control scheme, and its use in a single channel distributed control access protocol environment. The concept is simulated with the same three "Listen Before Talk" class of protocols used in Doc P802.11-92/51. The medium re-use potential is explained, and the relevant requirements for the MAC/PHY interface are given

Conclusion: Medium re-use efficiency is a major concern for especially single channel systems. As explained in Doc P802.11-92/51, the medium will be (co-channel) interference limited rather than noise limited. A method to improve the medium re-use efficiency has been introduced, that dynamically limits the Transmit power (and an associated "Defer" threshold) to a level needed for reliable communication per individual destination.

Analyses and simulations presented have shown that large re-use efficiency improvements can be obtained depending on the physical size of the BSA. This Dynamic Power Control facility can be an option, and it can work in a mixed environment.
The only thing we need to standardize on is two extra Bytes in the MAC frame structure that contain settings for the transmit power with which the packet was transmitted, and a "Silence Level" Byte that represents the average interference level at the location of the transmitter. The resulting requirements for the MAC/PHY interface would be to have a MAC to PHY management interface available for control and status exchange, over which the MAC should have full control.

Introduction:

Contribution P802.11-92/51 compared different distributed access protocol alternatives. One of the relations simulated was the Total Throughput as function of the distance between two networks in both a peer-to-peer (Fig 6), and in a Client-Server (Fig 9) network load. Individual station behavior was shown in Fig 7 and 10 respectively. It should be noted, that this simulation involved a homogeneous environment, with the average attenuation coefficient that represent a typical open office environment. Although the results show the sharing effects, it is less realistic for practical use. In reality you won't find buildings that are 800 meters square, without any major attenuation obstruction like concrete Walls, and things like elevator shafts, which highly effect the propagation conditions over larger distances.
Therefore it is more appropriate to look at Fig 7a and 10a where an additional 20 dB attenuation step is assumed between 2 networks, which can be considered as an effect of a concrete wall, or floor to floor attenuation. This still means however that under these conditions and with the sensitivity levels used, the same medium must be shared between a considerable number of BSA's. It looks like that under the described circumstances the medium needs to be shared with 4-6 other networks on the same floor, and at least one on the floor below, and one floor above for a multi-story building.

To assure proper receiving conditions, a sufficient Signal to Interference Ratio (SIR) is required at the receiver. Therefore it is required to force other medium contenders to be silent in an area around the intended receiver. The area that needs to be forced silent will depend on the received signal level of the wanted signal at the receiver. When the transmit power is adjusted to a level suitable to assure a required SIR, significant lower power levels can be used for all stations closer to the transmitter.
This will in turn reduce the co-channel interference levels in neighbor BSA's, and consequently allow for a better re-use of the medium.

Please note that this is completely independent of the access protocol used. This paper proposes a method that allows for Dynamic Power Control, and introduces the functions required in the MAC and PHY, to support this. Therefore it presents the functionality needed in the MAC/PHY interface.

Power Control System description

In the proposed system the transmit level, and the defer threshold can be changed as a function of the destination of the packet. This means that for close-by destinations much less transmit power can be used to achieve an acceptable SIR condition at the receiver station. Transmitters in nearby other networks can then utilize the medium at the same time, because they don't have to defer for the much weaker signal. This makes re-use of a certain frequency band possible.
When a transmitter wants to transmit to a nearby station with a reduced power level, then its defer threshold should be reduced as well to take advantage of the fact that the defer distance can be reduced. Otherwise this transmitter would still defer for any station within the maximum defer boundary (400 meter as illustrated in figure 1 and 2 in the example system).

How to determine the transmit and defer levels:

The idea is that addressed receivers can measure the signal level of the received signal coming from a specific source.
By default the attenuation between two points is unknown after startup of the system, so they should transmit at full power. Per connection the receive level should be averaged over a number of measurements to decrease the effect of fading when stations are moving.
For this method to work, the receivers should know the power level with which the packet was transmitted. This information can be included as a separate byte in the MAC frame structure. By reading this byte and subtracting the actual receive level, the signal path attenuation (at that moment) can be calculated. As already explained, by averaging of this attenuation value over a number of packets received from that source, the average path attenuation can be determined. This value can be an input to an algorithm that determines the transmit level and the defer threshold, using appropriate margins for fading. The resulting values can so be maintained by the MAC in a table, referenced by the destination address, so that every time a packet is to be transmitted, the relevant Tx-power level and defer threshold can be programmed from that table.

Note that different algorithms may be needed for mobile stations, than for fixed stations, because for mobile stations, the average attenuation will vary over time, so that the averaging algorithm time constant should be small to allow fast adaption to the attenuation situation.

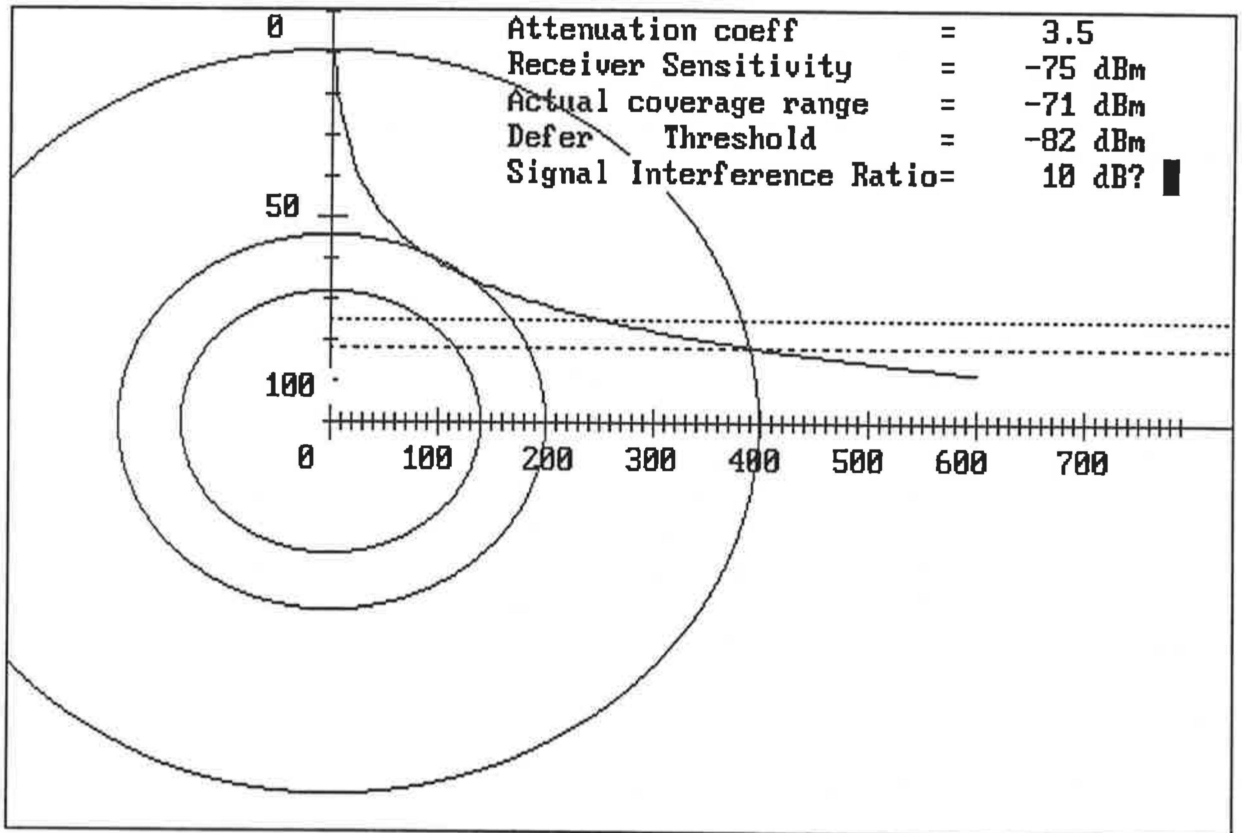

Figure 1

Explanation of Fig-1:

The figure shows two relations. One is the average probability receive signal level as function of distance to the transmitter (located at the origin). The transmit output level measured at 1 meter from the source is in this case -6 dBm. The shown average probability attenuation as function of distance is typical for an indoor open office environment.
For ease of explanation it is assumed that the environment is homogeneous, so not further separated by attenuation boundaries like walls or ceilings. In addition no margins are included yet to limit the outage due to interference.
Next on the same axis a number of circles are drawn. The outer circle represents the boundary at which the average attenuation of the signal is such that its level is equal to the "carrier detect" threshold of the receiver. The transmitter is assumed in the origin and the destination receiver

on the circle.
When the radio modems are using a CSMA type access protocol, then this is also the average boundary for which stations will defer for other transmitters in this area. Effectively all the transmitters in this area will share the same bandwidth. The middle circle would be the average boundary on which packets can be reliably received, when assuming a certain background noise level (in this case 18 dB above thermal noise).
However in order to operate with a certain bit error rate, the interference level from other network activity should be an amount of SIR dB's lower then the wanted signal. The SIR value depends on the modulation method and quality of the receiver implementation.
Now when we assume that just outside the outer circle there can be a transmitter active, because we will not defer for such a low level signal. Then all receivers within the inner circle will receive the wanted signal at least SIR dB higher than the disturbing signal. Therefore the inner circle represents the so called "Safe Coverage Area", which is typically considerably smaller than the middle circle representing the boundary for noise.
On top of this a fading margin needs to be included, to assure a low enough outage probability due to fading.

Advantage of dynamic Power Control:

When the destination of the signal is for instance located at half the "Safe Coverage" distance, then in case of a fixed power level situation the wanted signal is much larger than the interference signal, which will assure reliable operation.
However, because of the defer mechanism, no other transmission can take place at the same time in an area as large as the outer circle. When the transmit level would have been reduced to a level that is sufficient to guarantee a SIR difference between the wanted and interfering signal, then the interference level generated in other nearby networks will be lower, and consequently a transmitter much more closer by could be allowed to transmit.
At the same time however the "Carrier threshold" or "Defer threshold", of the station that wants to transmit at the reduced power level, should be reduced as well because otherwise we would unnecessary defer to nearby transmissions in this case.

It can be calculated that the overall re-use of the medium is approx. a factor 3 improved (for the homogeneous uniform station distribution case) when dynamic power control is used.
An additional major advantage is that a dynamic system like this will adjust its transmit power level, and associated carrier threshold such that small clustered networks will generate much less interference to neighbor possibly larger networks, so that the total throughput of the system will significantly increase.

When fixed levels are used, then the transmit level and receiver thresholds are to be set as a trade-off between coverage distance and throughput/sharing, depending on the application. This will however also mean that when we have only a relative small network clustered over a small distance then we still use the full power and receiver sensitivity, so that we have to share the medium with all neighboring systems within a pretty large area, although for reliable operation we could do with a lot less power. Also the relative throughput of the system will increase significantly when only enough power is used as needed for a small clustered network like described. Given this situation, then the re-use improvement can be much higher than a factor 3, and may easily be in the order of 30 or more.

How to determine the Tx-level and Rx threshold?

The attenuation between transmitter and receiver can be determined in the receiver by measuring the received level in the receiver, and subtracting the transmit level. This transmit level should be indicated in a byte in the MAC frame. In addition to that also the "silence" level as measured at the transmitter side should be included as information in the MAC frame. At a receiver location this calculated attenuation and the silence level should be stored in a table indexed by the address of the transmitter. Given the measured attenuation, an optimum transmit level can be calculated, or looked up in a pre-calculated table, together with a new "defer" threshold. The "silence" level should be used to verify that the expected receive level is at least SIR dB higher than this "silence" level. If not, than the transmit level should be increased accordingly. Note that the "silence" level is a measure for the continuous interference level at that particular receiver location.

How to learn the system.

The system should start at default maximum transmit levels and associated defer thresholds, and will learn the more optimum settings when communication progresses. No interchange of information other than the silence and power level included in the MAC frame are needed.
The system could learn from only the communication addressed to that workstation, and build a table only for those stations that it communicates with.
It could also learn from every packet crossing the medium, and maintain a table for all communication it can hear. This will cost large tables, also for those destinations which will probably never be addressed by this particular workstation, so this method is not recommended.
Usually it is important to communicate with a server station. When packets from server stations or access points to a wired backbone are uniquely identifiable then the system could learn from every transmission from such a server station.

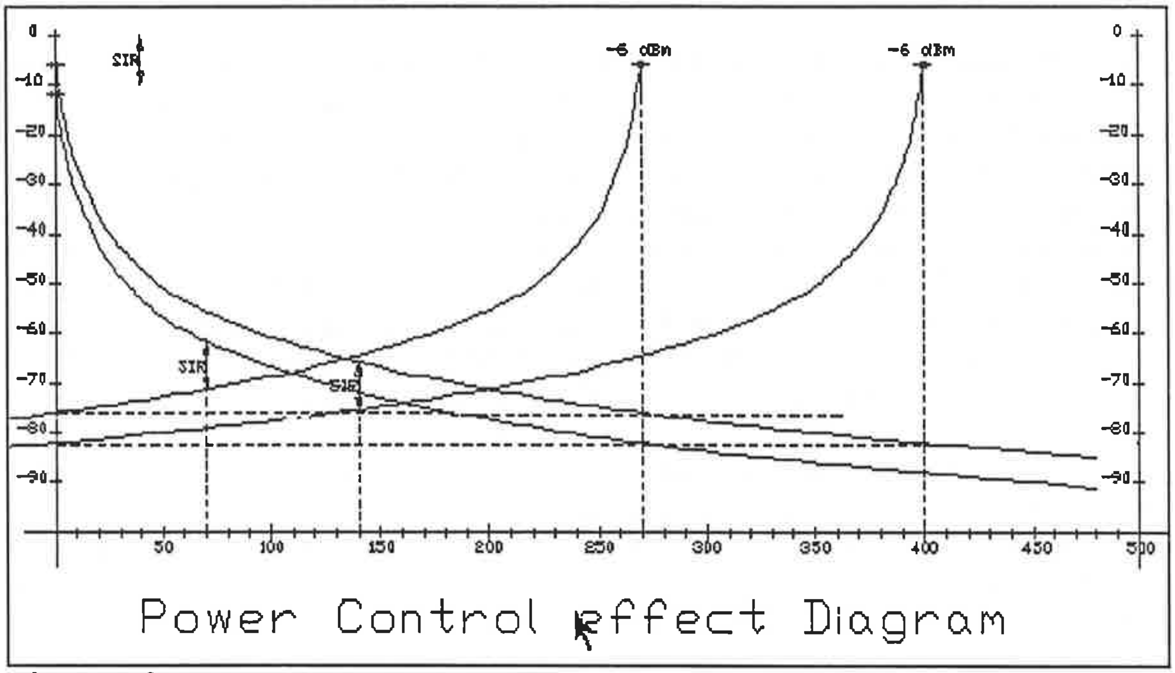

Figure 2

Note: For ease of explanation, only the average probability curve is shown. In practice an appropriate fading margin needs to be added (to the SIR).

Explanation Figure-2:

In figure 2 again a level versus distance plot is given for a typical office environment, but the environment is assumed homogenous without walls. It shows that when the Defer threshold is -82 dBm, the transmitter will defer for all transmitters within a 400 meter range. When an interfering station is active just outside this boundary then all stations within a 140 meter radius from the transmitter will meet the SIR criterium.
Now assume that an intended receiver is located at 70 meter from the transmitter, then this location could be reached with a 6 dB lower transmit level, while sufficient SIR separation is assured for the nominal level transmitters located as close as 270 meter. Also the defer threshold can be decreased with 6 dB to allow start of transmission when no other transmitter is active within 270 meter rather than the original 400 meter.

Potential medium re-use gain.

It has been analyzed that the proposed method has the potential to increase the medium re-use efficiency by a factor 3 in a homogeneous uniform station density environment.
In practice this can be much larger and is a function of the size of the network.
Simulations have been done using the simulator describe in Doc P802.11-92/26, in which the power control mechanism is implemented. Some result plots are shown in the appendix. The plots show the relation between total throughput versus the distance between two networks, for both a peer-to-peer traffic model as well as a Client-Server traffic model. To calculate the re-use gain, it is a good practice to take a point on the throughput curve which lies in the middle between the lower and higher boundary, so the 150% point (150% of the capacity for one NW). Re-use efficiency can be expressed as the increase of the capacity per unit of area. So the improvement would be the relation between the square of the distance at the 150% point of the two curves. So this is the ratio of the area's at those points. This ratio is called the Power Control Efficiency (PCE) value, and is calculated for the plots shown.
The simulation results show typical numbers for the Client-Server traffic model used on a network layout as shown in doc P802.11-92/51. Also results are shown for the peer-to-peer traffic model, which show a lower efficiency gain.
This can be explained by the fact that the distances in the peer-to-peer case are larger, so less power reduction can be obtained. This is equivalent to a double size Client-Server based traffic model with the server located in the center of the area. Also the fading margin used during the test was too low (only 5 dB), so that more retries are needed.
Additional plots show the Client Server results for a smaller size network, which show higher PCE numbers as expected. Note that for these plots a more realistic 10 dB fading margin is assumed, so that the plots do not directly compare.

The following test conditions were used for the simulations:

Client-Server traffic model:	RW (Read/Write) mode with 60% Short 64 Byte packets and 40% Long 512+64=576 Bytes packets. Performance results show 512 Byte data only throughput, so without Novell overhead.
Peer-to-Peer traffic model:	RW (Read/Write) mode with 60% Short 64 Byte packets and 40% Long 512+64=576 Bytes packets. Performance results show all MAC throughput data only throughput.

Fading margin used in Power Control algorithm is only 5 dB unless otherwise indicated. Note that this is too low for an acceptable outage.

Power Control as an implementation option.

Stations could implement Power Control as an option. That means that a mixed environment could exist with one BSS using Power Control, while a partly overlapping other BSS would not use it. Even within a BSS, there could be a mix of stations with and without the Power control feature. All these mixes are possible as long as the extra Power Control Bytes in the frame structure represent the fixed settings of the transceiver. This is because the Power Control algorithm is distributed over the stations, and only needs to operate in the stations that have this option installed. Of course the overall throughput improvement would be limited, but stations that are relative close to their main destination (AP or Server), would benefit from having this option installed.
This characteristic is important, because it allows for optional lower cost PHY implementation, because those PHY's could be implemented without a variable gain powerstage, and without a receive level measurement facility.
The only bottleneck would be the "Silence" level. When a station with power control communicates with a station or AP without any power control provisions, then when the destination is located near a CW type interferer, then the power control would hurt, because the "Silence" or interference level is not known. A station could learn that something is wrong, by maintaining retry statistics, which could cause the Power Control to be turned off for that particular destination.

Power Control in a multi channel system.

When sufficient channels are available, together with an efficient allocation plan or algorithm, to assure that there is sufficient re-use distance, then Power Control would not be needed, or would have only minor effect.
This is when we assume that only one ring of cells with different frequencies (6-10) is needed for isolation, because it is likely that also a wall will be in between the frequency re-use area.
This will of course depend on the size of a BSA. When the range of a single BSA is very limited to for instance 10-20 meter, then it is likely that a high density is needed to cover one room. So that re-use is needed within one room, or additional cell rings are needed for isolation, requiring more channels (approx 20). With Power Control, re-use can be improved so that less channels are needed for re-use isolation.
Of course the standard should not require Power Control at all times, because there will likely be multiple channel PHY's, or low requirement implementations that don't need Power Control.

What provisions are needed:

The following functionality is needed in different area's to support dynamic power control:

- In the PHY:
  - A variable gain power stage
  - A received signal level measurement facility
  - A variable "Defer" threshold

- In the Frame:
  - A byte containing the level with which the packet was transmitted.
  - A Byte containing the averaged "Silence" level measured at the transmitter location.

- In the MAC:
  - Facilities to dynamically control the Tx-level, and "Defer" threshold per packet.
  - Facilities to retrieve the Receive level and "Silence" level from the PHY.
  - Power Control Learning algorithm.
  - A table with Tx-level and Defer threshold level entries per destination.

It may be considered a disadvantage that the MAC needs to maintain a table per destination, with the TX-level and Defer Threshold level. This is likely a software function, and may burden the host system. The advantage can however be that the same table can be used for other purposes. It makes it for instance possible that the bitrate can be changed dynamically per destination station. This would allow to operate most of a network at a high datarate, while perhaps a few far away stations suffering from much path attenuation, or from a difficult delay spread situation can be served at a reduced bitrate.
Given that AP's need to maintain similar tables anyway, (for bridging, power management and others), then this would have hardly any effect on the complexity.
Stations will in practice communicate with only a limited set of destination stations, of which most could be located beyond the AP on the backbone, so only a limited table needs to be maintained.

MAC / PHY Interface requirements:

Figure 3 shows a functional block diagram, indicating the additional functionality needed in the interface between the MAC and the PHY. Figure 4 shows a hardware implementation level blockdiagram in which the extra functions needed for the support of Power Control are shown. The additional algorithmic requirements needed are not shown, but these will likely be software functions needed in the Driver Software.

The MAC will need to manage the PHY on a per packet basis, and will either need a dedicated management interface for this, or will need control over any existing management interface. Before the medium is accessed, the Tx-PowerL and Rx-THRhold values need to be downloaded in the PHY, after which the MBUSY interface signal is being sensed as part of the "Listen Before Talk" class of distributed access protocol. At the end of a received packet, the Rx-level and Silence-Lvl need to be retrieved from the PHY, to attach it to the received packet in its receive frame buffer. The "Rx-Level" can be measured anywhere within the received packet. The "Silence-Lvl" should be measured in a defined silence period such as the Inter Frame Space (IFS).
It is likely that at this moment also other PHY status is collected that also need to be attached to the received packet. Examples will be any Quality Of Service (QOS) related status like "signal Quality".
Perhaps the TxD and RxD data lines can be multiplexed for this, to serve as a management datapath.

More research needed:

More research is needed to analyze the required specifications, and to develop a suitable learning algorithm. For instance the effect of the tolerance of the Transmit output level and the receive level measurement unit must be evaluated. Further analyses is needed to specify the dynamic range and discrete step sizes needed for the variable transmit gain stage. Possibly a different parameterized learning algorithm is needed for mobile stations.
The conversion of measured level to desired output power is an other field of research.

Conclusion:

A method to improve the medium re-use efficiency has been introduced, that dynamically limits the Transmit power (and an associated "Defer" threshold) to a level needed for reliable communication per individual destination.
Analyses and simulations presented have shown that large re-use efficiency improvements can be obtained depending on the physical size of the BSA. This Dynamic Power Control facility can be an option, and it can work in a mixed environment.
The only thing we need to standardize on is two extra Bytes in the MAC frame structure that contain settings for the transmit power with which the packet was transmitted, and a "Silence Level" Byte that represents the average interference level at the location of the transmitter. The resulting requirements for the MAC/PHY interface would be to have a MAC to PHY management interface available for control and status exchange, over which the MAC should have full control.

DYNAMIC POWER CONTROL POTENTIAL

* Problem addressed

  - Environment is INTERFERENCE LIMITED

  - In single channel environment, medium needs to be shared with a significant number of different BSA's.

  - In multi channel environment a large number of channels may be needed for re-use isolation.

  - For large cells, need 6-10 channels.

  - For small cells, need around 20 channels.

* Proposed solution:

  - Reduce the co-channel interference.

  - Limit Tx Power level to a value needed for reliable reception.

  - Learn proper level per individual Destination.

  - Use Defer threshold as function of Power Level.

DYNAMIC POWER CONTROL POTENTIAL

* Features

  - Analyses show factor 3 Re-use efficiency improvement in a homogeneous uniform distributed environment (Assuming single channel system).

  - Larger Re-use efficiency as function of BSA size.

  - Could be protocol type independent.

  - Can operate in mixed environment. So can be implemented as an option.

  - Easy extension to mixed bitrate environment.

* Simulation Results:

  - Simulator described in P802.11-92/26 includes Power Control option.

  - Simulated with CSMA/CA, CSMA/CA + Ack, and 4-WAY LBT.

  - Results show Network size dependency

  - Good results for all 3 protocols.

DYNAMIC POWER CONTROL POTENTIAL

* Provisions needed

  - MAC Frame structure: (Needed in Standard)

    . One Byte indicating Tx-Power level
    . One Byte indicating "Silence" / Interference Level

  - PHY:

    . Variable Gain Output stage
    . Receive Level Measurement function
    . Variable Defer threshold

  - MAC:

    . Dynamic Control of Tx-Level and Defer Threshold
    . Retrieve Receive level and "Silence" level
    . Learning algorithm
    . Maintain Tx-Level/Defer-THRhold table per destination

  - MAC / PHY Interface:

    . Management interface under MAC control

* Further research needed:

- Effect of hardware tolerances
- Dynamic range and step size requirements for Power Control
- Learning algorithm
- Static / mobile difference

* Conclusion:

- Dynamic Power Control has good medium re-use efficiency Potential
- Also applicable to multi channel systems
- Can work in mixed environment
- Can be an implementation option
- Only need Frame Structure provisions in standard
- Need MAC/PHY management interface under MAC control

Description of the following figures:

- Fig 3: Layered architecture Blockdiagram showing the extra interface functions needed for Power Control.

- Fig 4: Implementation level MAC + PHY Blockdiagram showing the extra functionality required for Power Control. Also the required entries in the MAC frame structure are indicated.

- Fig 5: Shows the attenuation versus distance diagram, with various curves explaining the Power Control meachanism.
  It also shows the MAC frame structure with the two Bytes TPL and SSLVL required for dynamic Power Control. The sequence and place of the MAC frame structure is not relevant.

- Fig 6: Shows the Potential Power Control efficiency for the network used in doc P802.11-92/51 with a peer-to-peer traffic model, with random traffic destination. The total system throughput of two networks is shown as a function of the distance between two networks.

- Fig 7: Shows the Potential Power Control efficiency for the same Network, but with a Client-Server traffic model. This means that all traffic is with a centrally located Server station. Throughput is actual Data throughput only, so without Novell overhead. Thsi is equivalent to a Novell Perform3 test.

- Fig 8: Shows the same but now with 20 dB extra attenuation between the two Networks.

- Fig 9: Shows the same relation as in Fig 8, but now with a smaller size Network, and the Fading margin is increased from 5 dB to 10 dB.

- Fig 10: Shows the same smaller Network but without the 20 dB attenuation.

- Fig 11: Shows the number of succesfull packets from the Server to the individual stations in the 5 seconds test time.

Note: All simulation results show a Potential Power Control comparison between CSMA/CA + Ack, 4-WAY LBT and WAVELAN CSMA/CA protocol. WAVELAN is however not equipped with a Power Control facility.

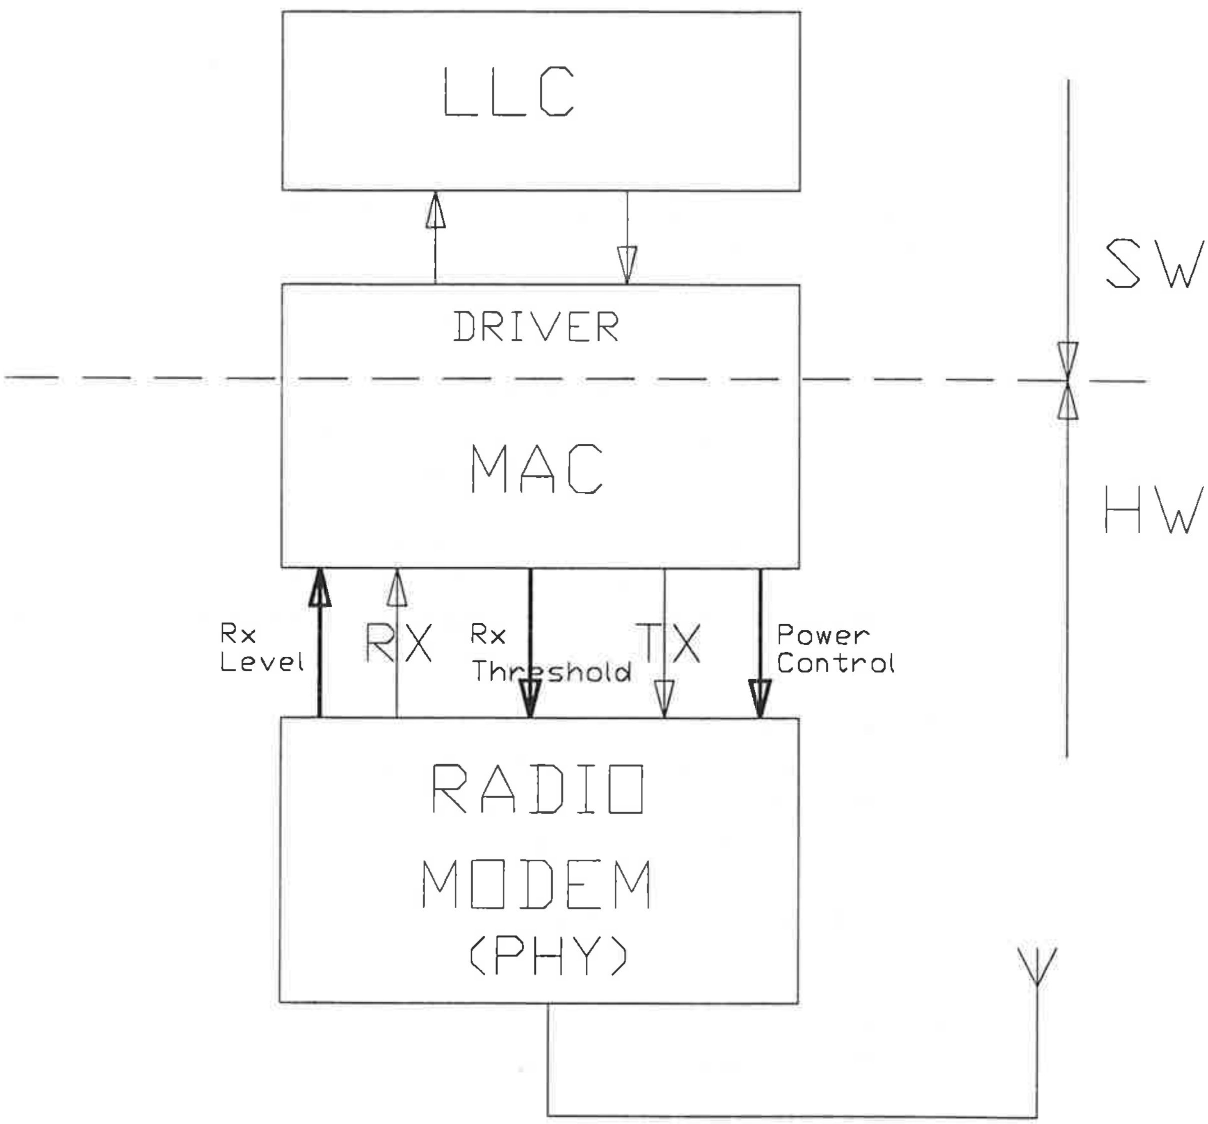

Fig. 3

WIRELESS COMMUNICATION CHANNEL
LAYERED ARCHITECTURE

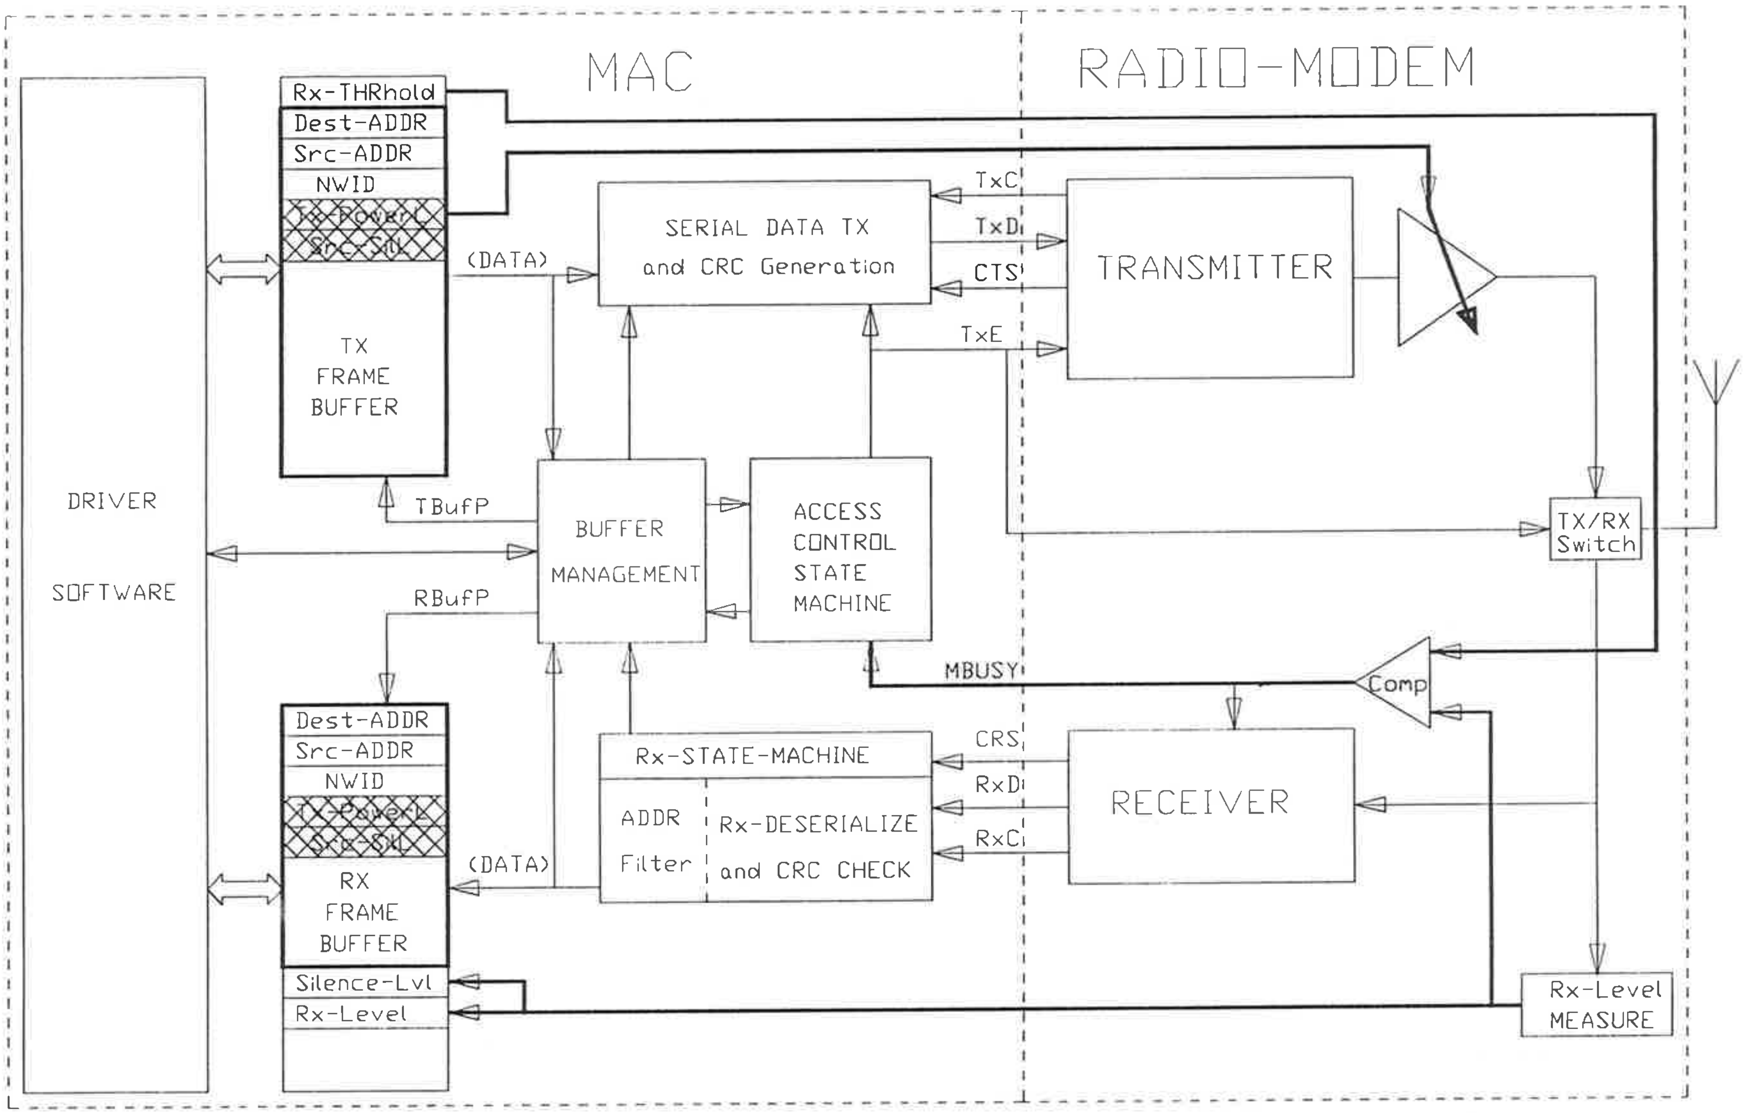

Fig. 4

WIRELESS MAC+PHYSICAL LAYER BLOCKDIAGRAM

MAC FRAME STRUCTURE

PREAMBLE	DEST ADDR	SRC ADDR	LNGTH	NWID	TPL	SSLVL	DATA	CRC
1	6	6	2	2	1	1		4

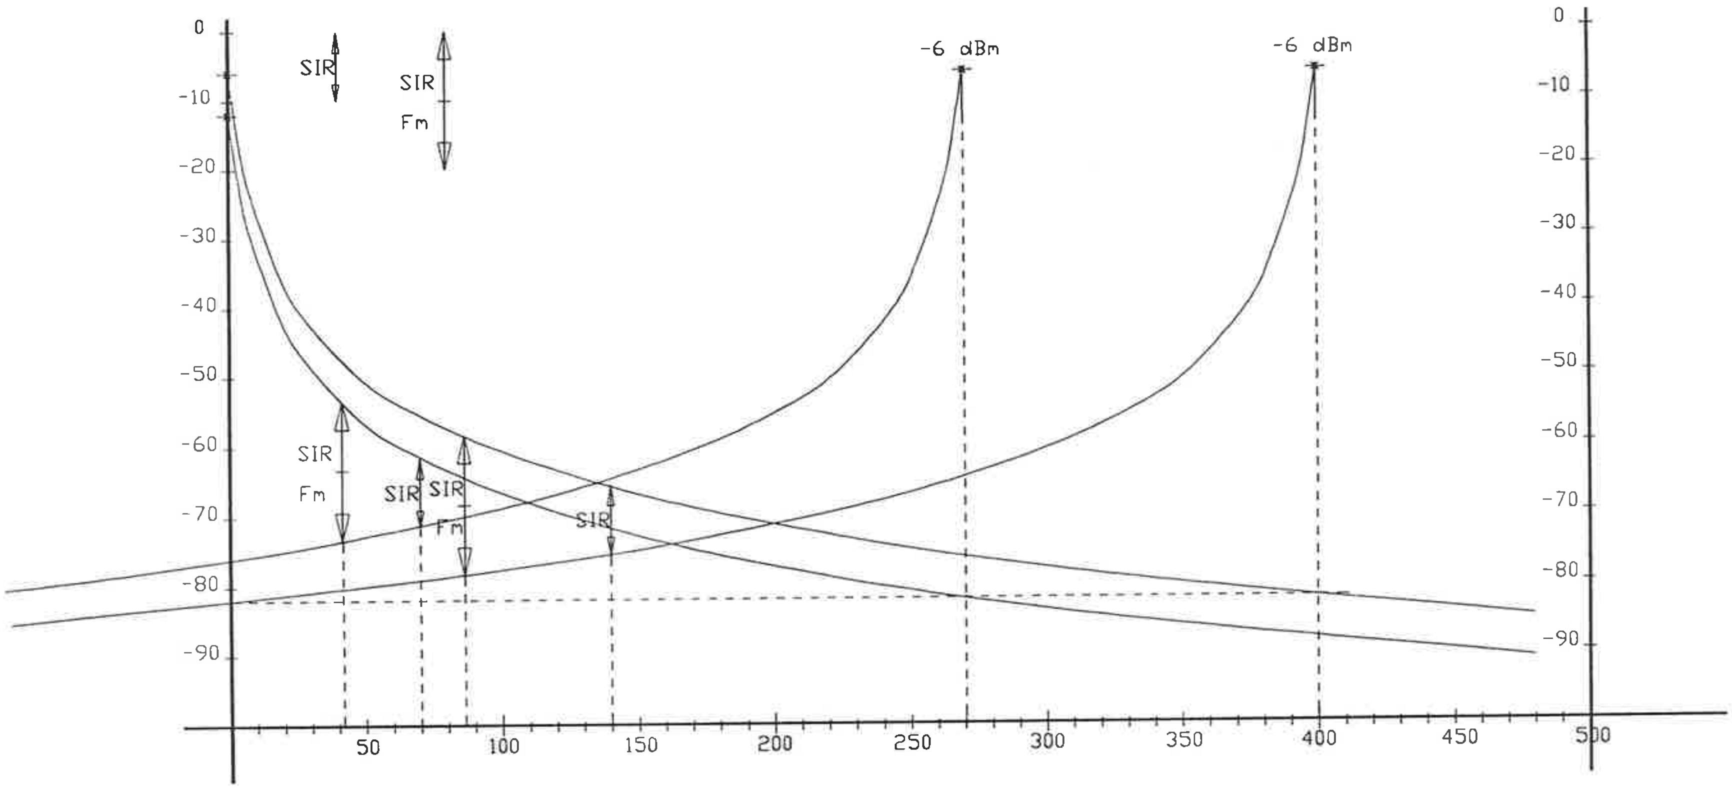

Fig: 5 Power Control effect Diagram

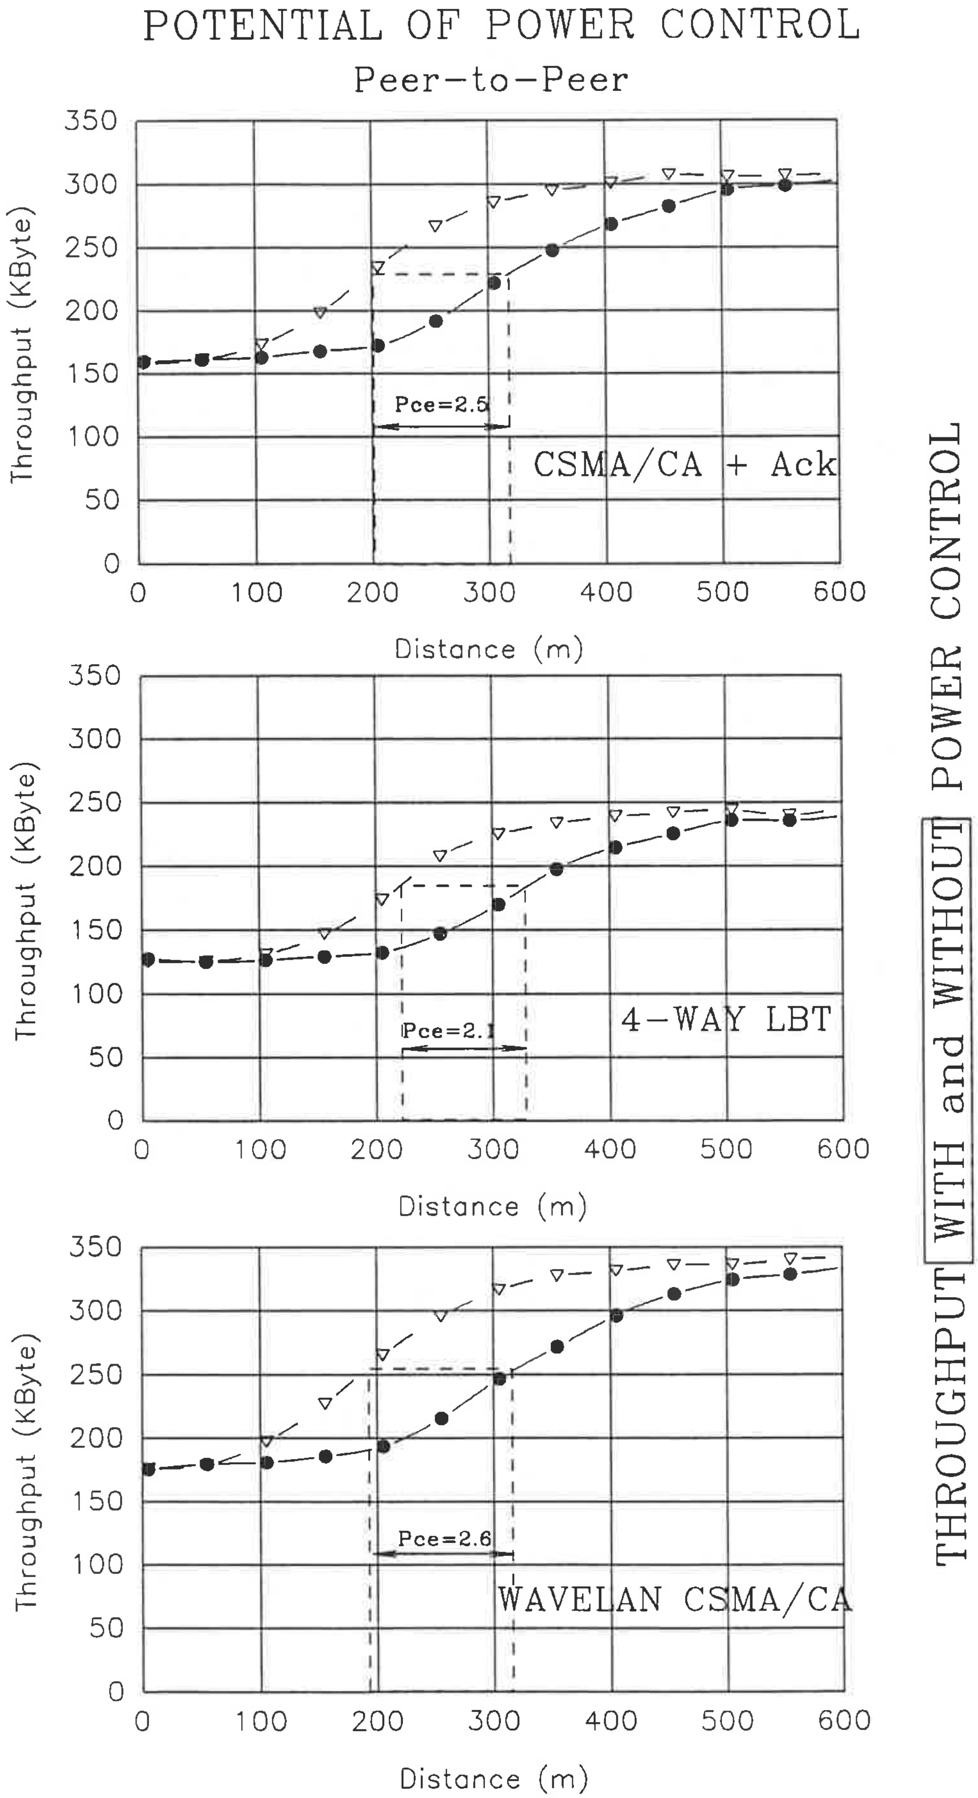
POTENTIAL OF POWER CONTROL
Peer-to-Peer
Throughput (KByte)
Distance (m)
Pce=2.5
CSMA/CA + Ack
Pce=2.1
4-WAY LBT
Pce=2.6
WAVELAN CSMA/CA
THROUGHPUT WITH and WITHOUT POWER CONTROL

Fig. 6

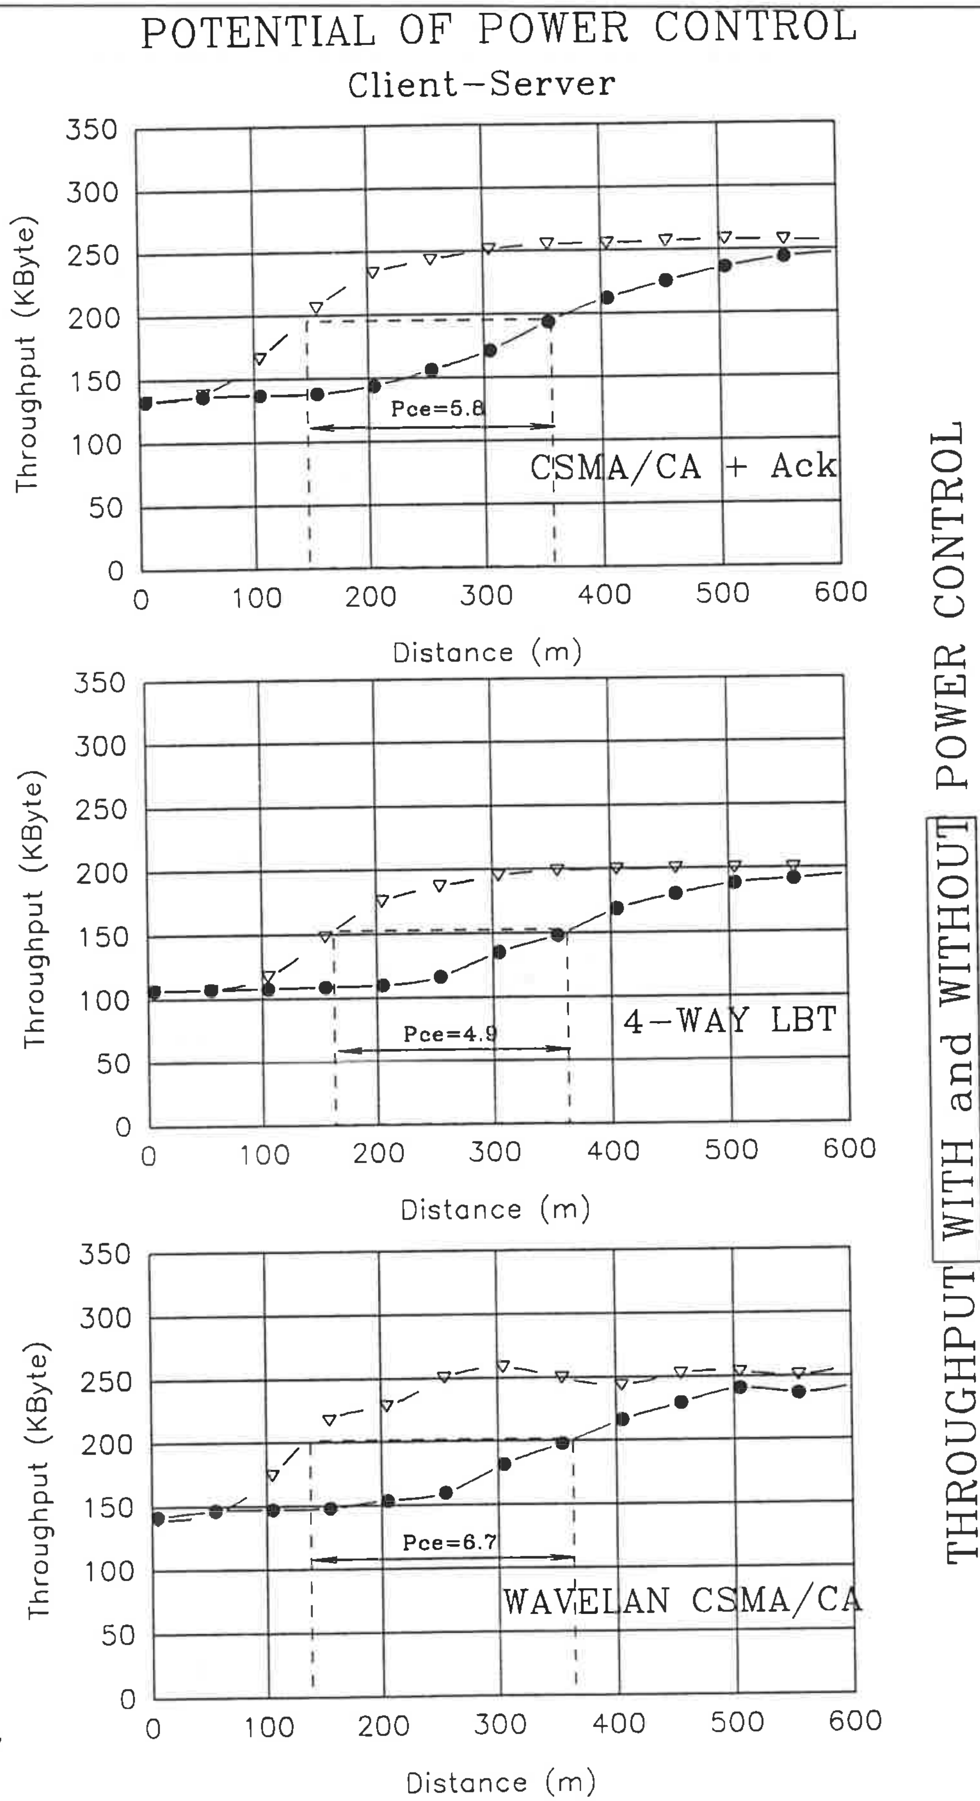
POTENTIAL OF POWER CONTROL
Client-Server
THROUGHPUT WITH and WITHOUT POWER CONTROL
Throughput (KByte)
Distance (m)
Pce=5.8
CSMA/CA + Ack
Pce=4.9
4-WAY LBT
Pce=6.7
WAVELAN CSMA/CA
0
50
100
150
200
250
300
350
400
500
600

Fig. 7

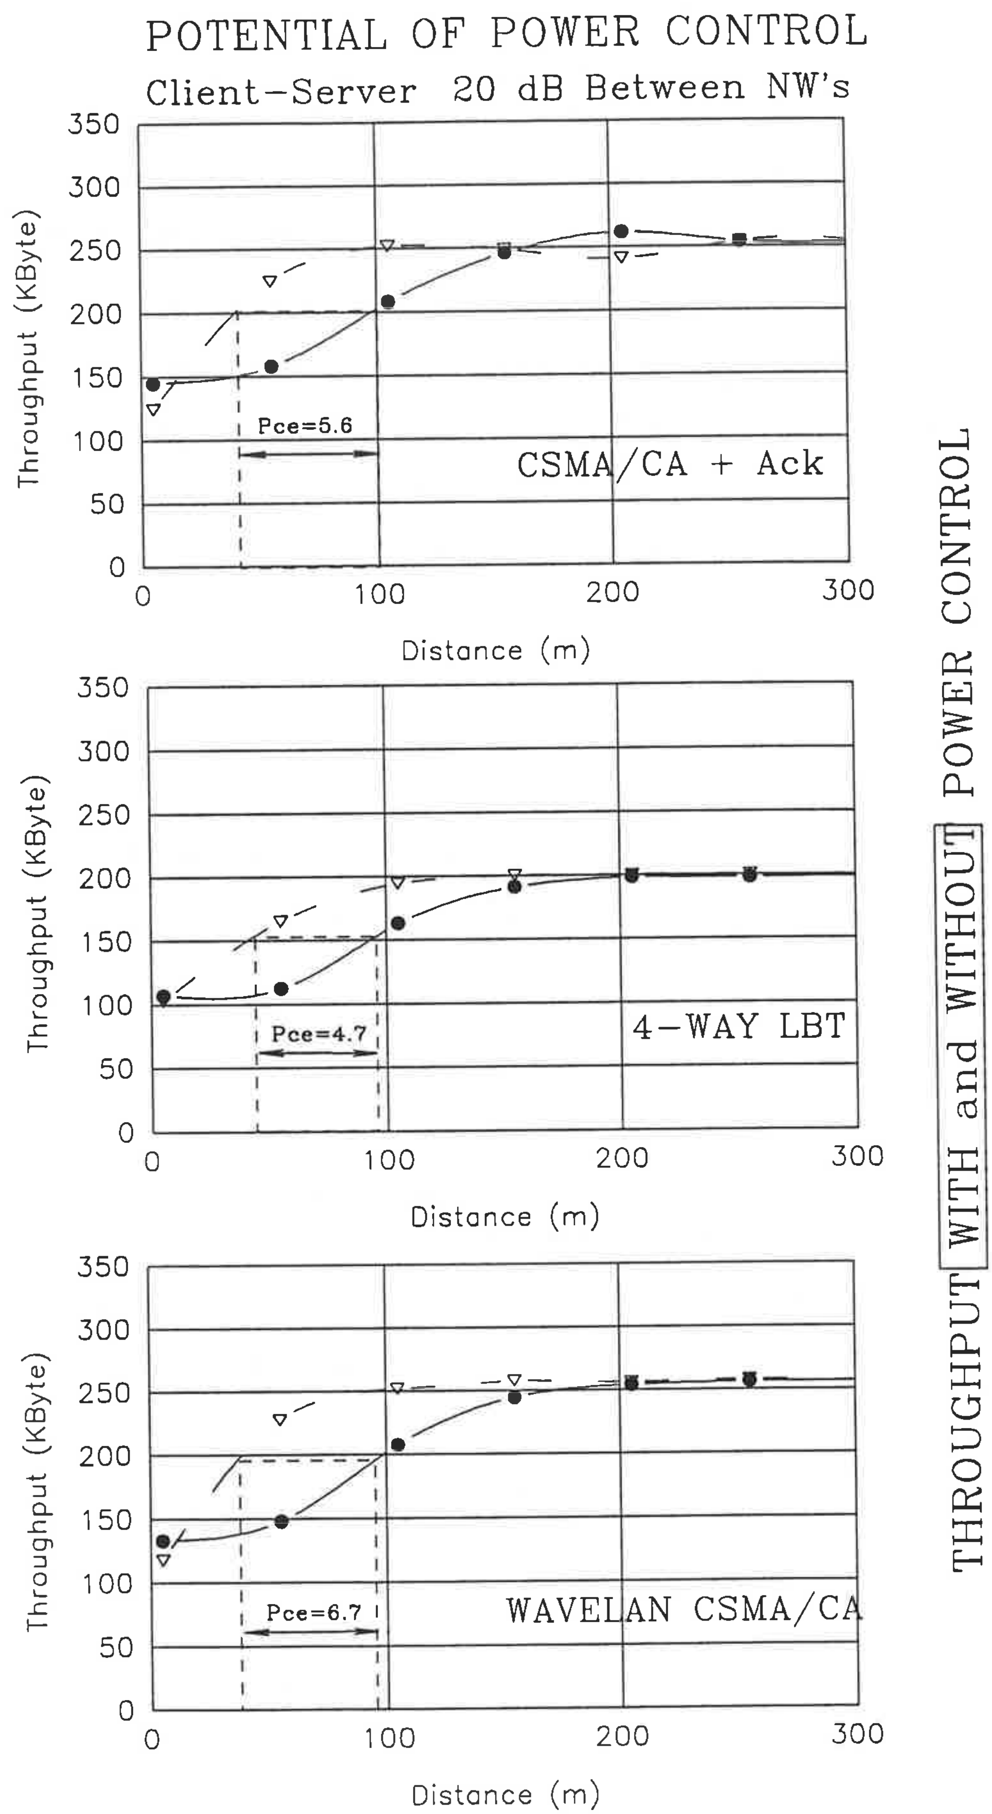
POTENTIAL OF POWER CONTROL
Client-Server 20 dB Between NW's
Throughput (KByte)
Distance (m)
Pce=5.6
CSMA/CA + Ack
Pce=4.7
4-WAY LBT
Pce=6.7
WAVELAN CSMA/CA
THROUGHPUT WITH and WITHOUT POWER CONTROL

Fig. 8

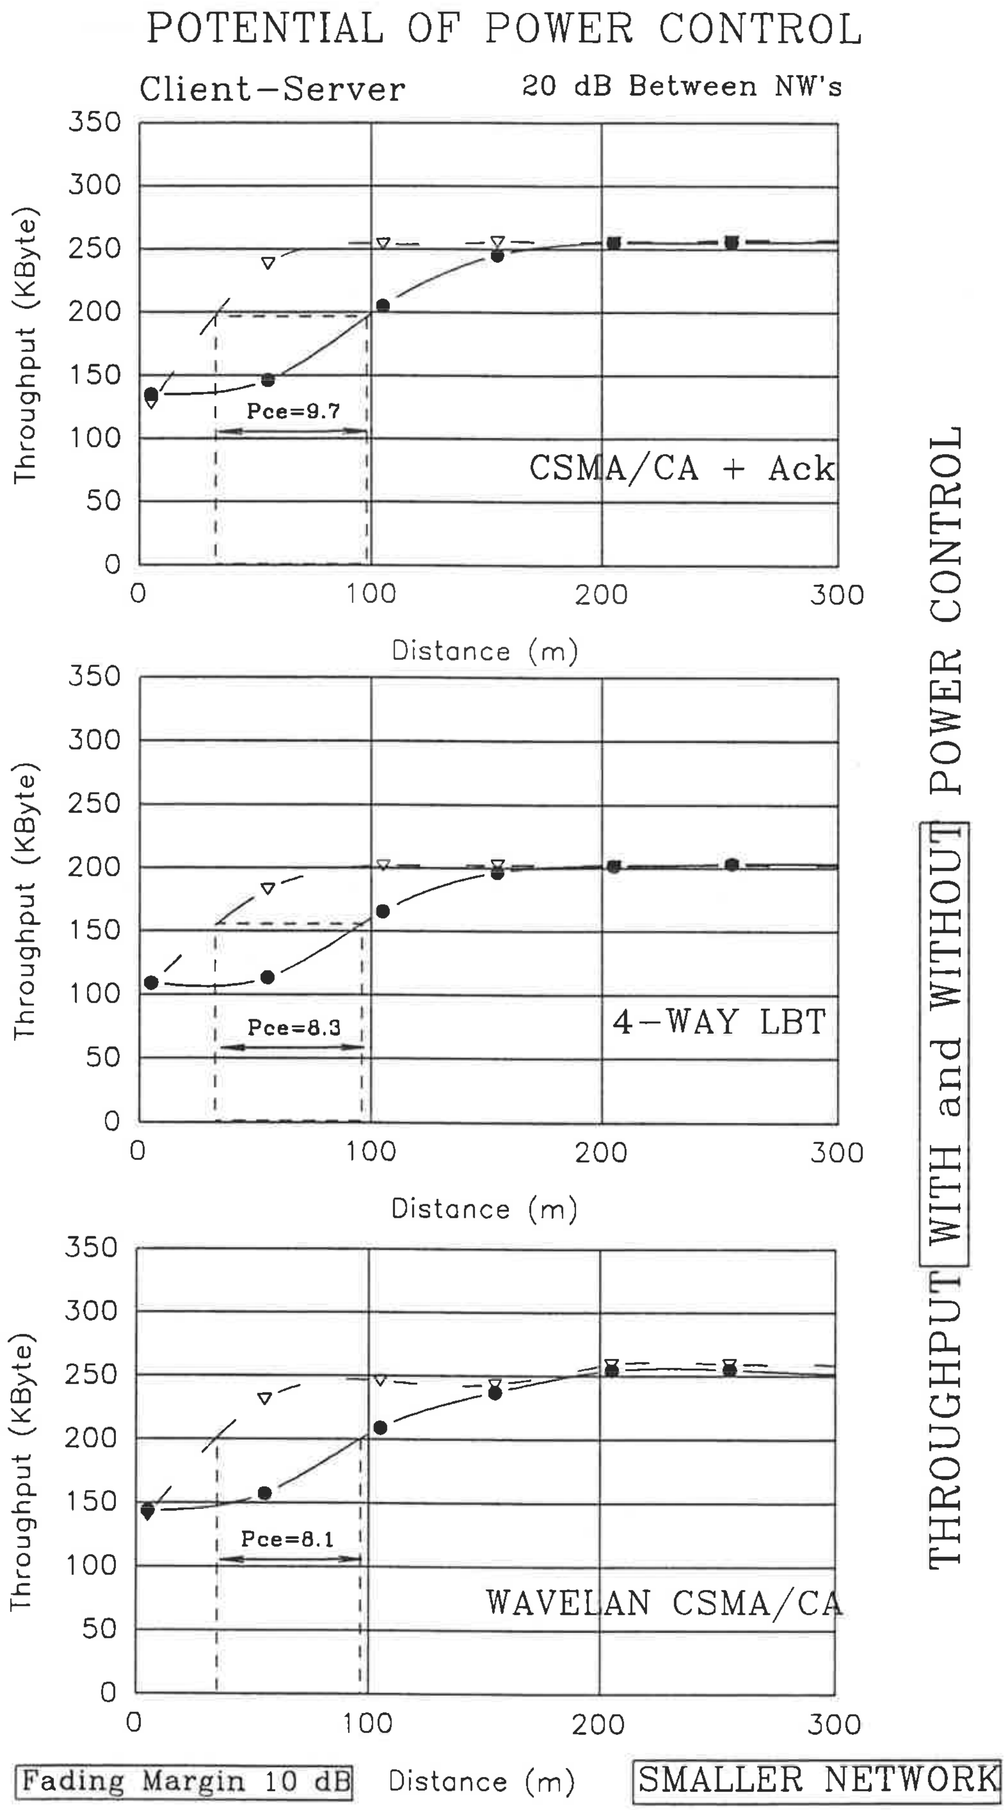
POTENTIAL OF POWER CONTROL
Client-Server
20 dB Between NW's
Throughput (KByte)
Distance (m)
Pce=9.7
CSMA/CA + Ack
Pce=8.3
4-WAY LBT
Pce=8.1
WAVELAN CSMA/CA
THROUGHPUT WITH and WITHOUT POWER CONTROL
Fading Margin 10 dB
SMALLER NETWORK

Fig. 9

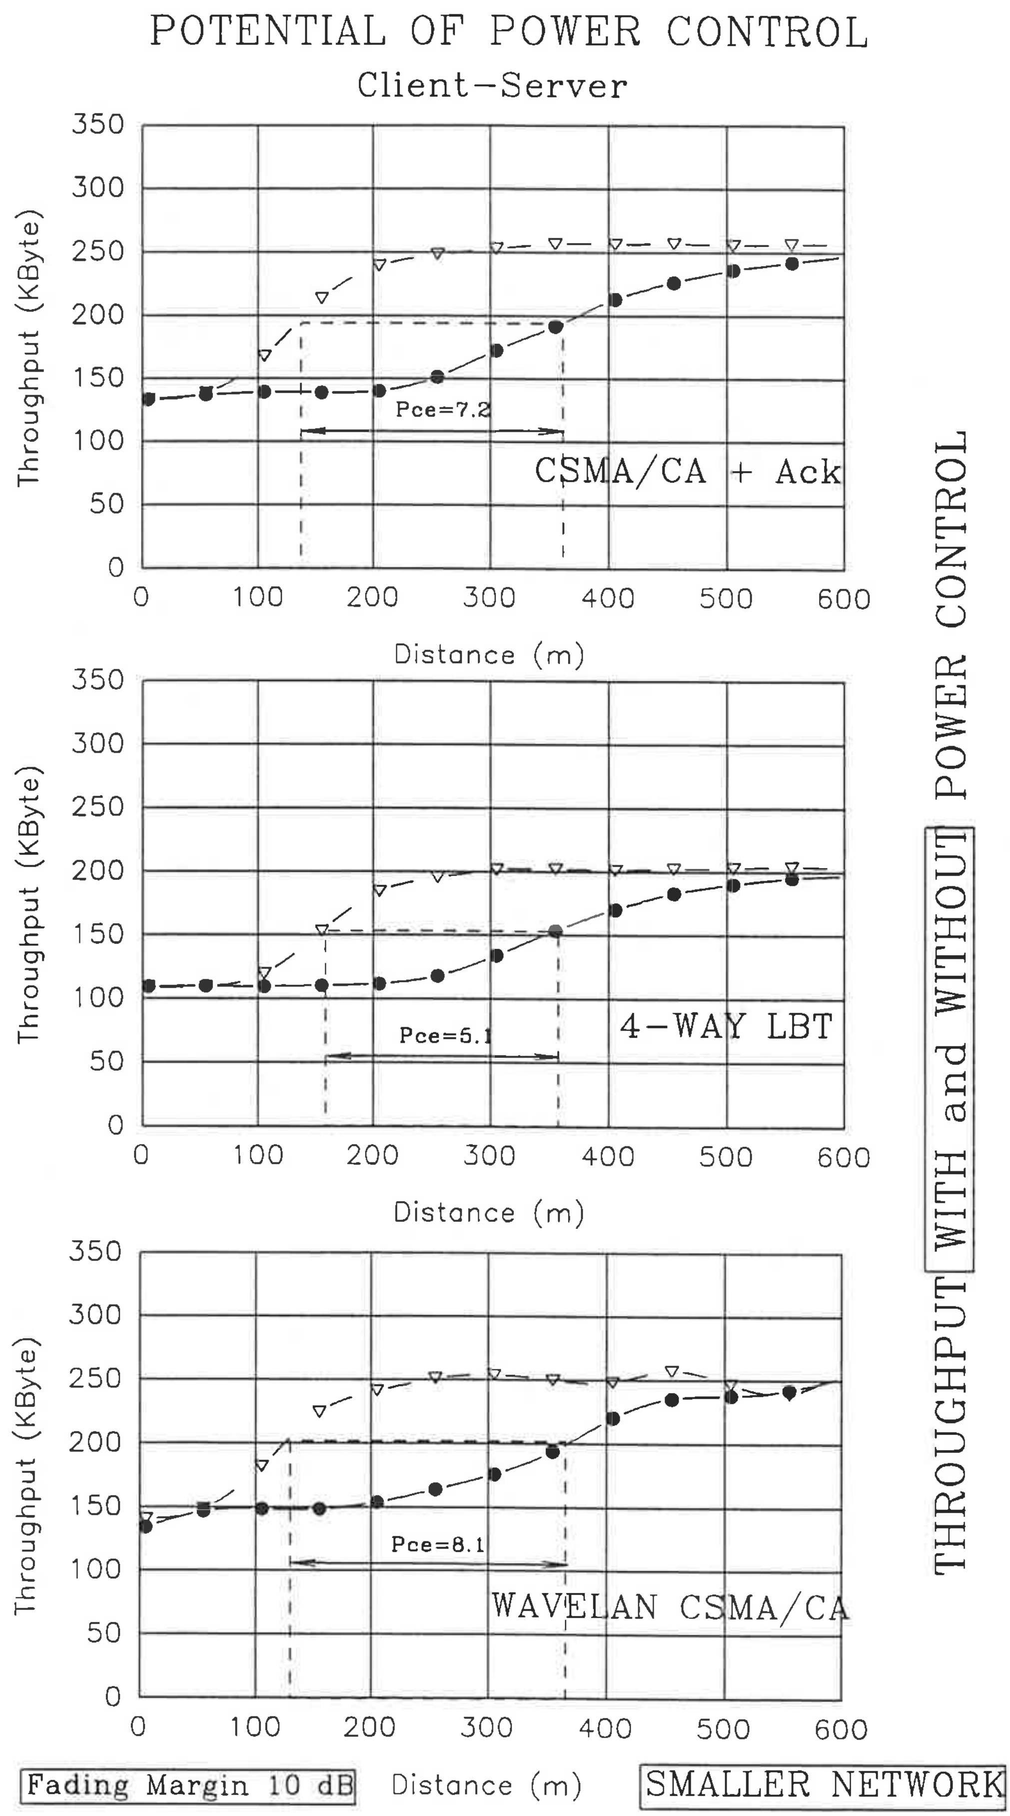
POTENTIAL OF POWER CONTROL
Client-Server
Throughput (KByte)
Distance (m)
Pce=7.2
CSMA/CA + Ack
Pce=5.1
4-WAY LBT
Pce=8.1
WAVELAN CSMA/CA
THROUGHPUT WITH and WITHOUT POWER CONTROL
Fading Margin 10 dB
SMALLER NETWORK

Fig. 10

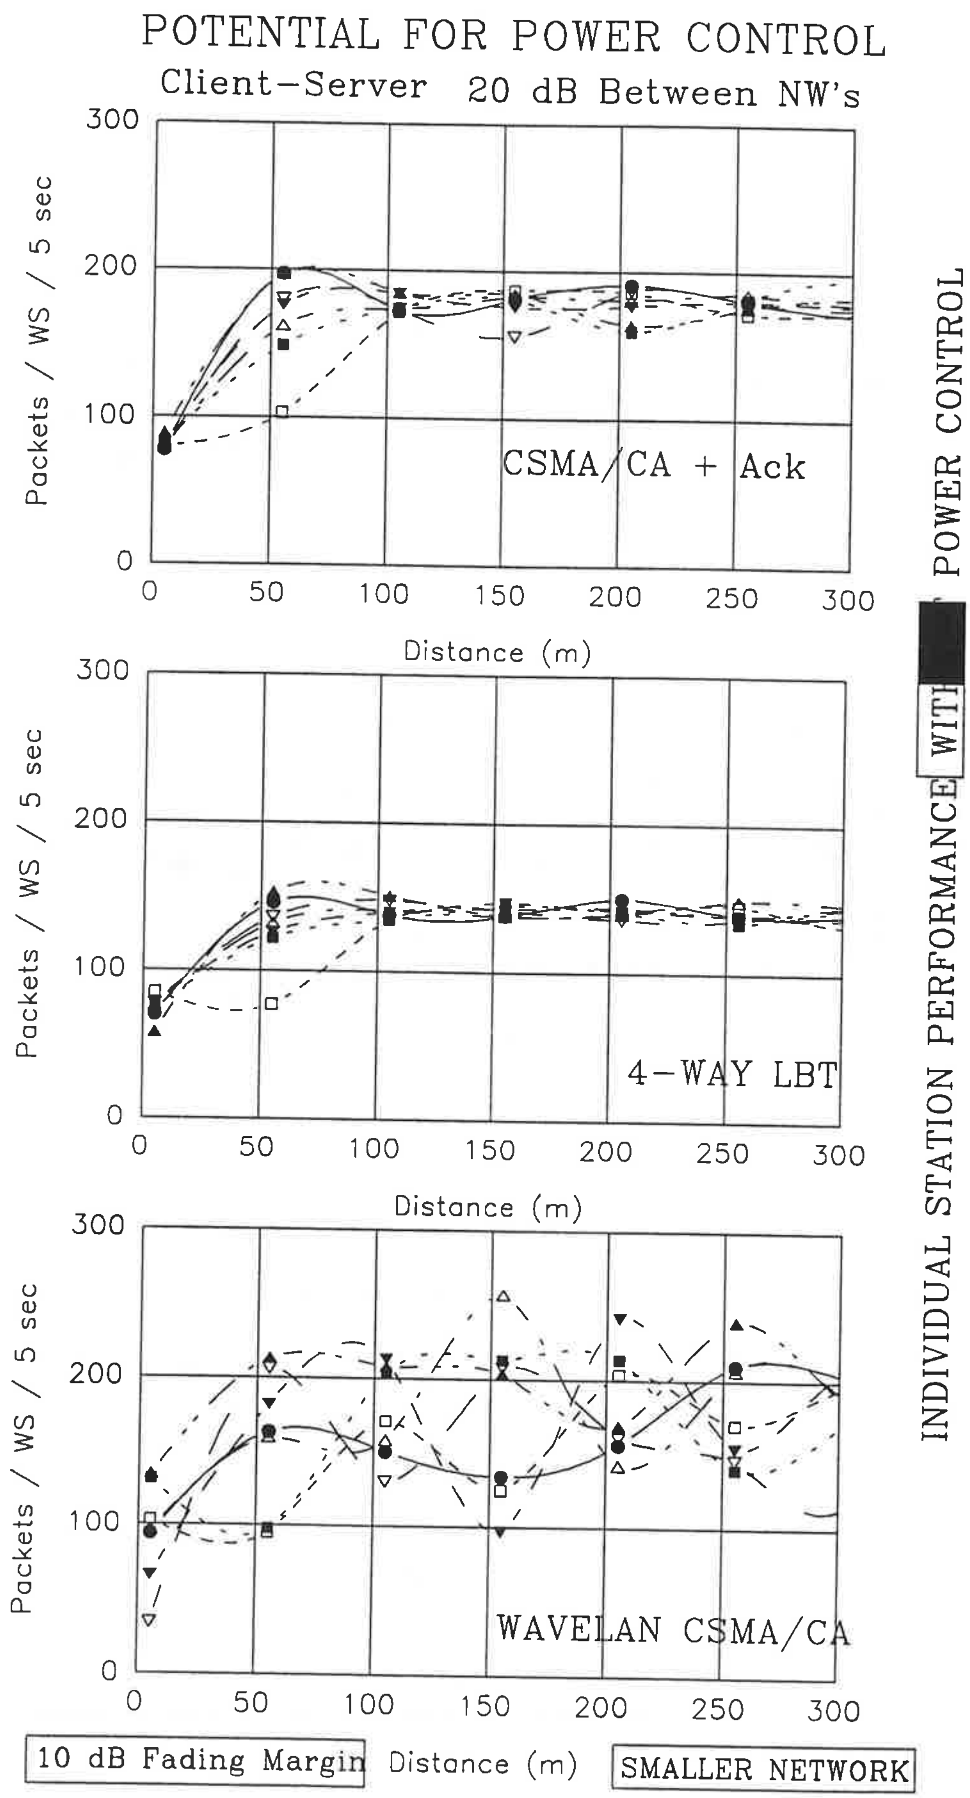
POTENTIAL FOR POWER CONTROL
Client-Server 20 dB Between NW's
Packets / WS / 5 sec
Distance (m)
CSMA/CA + Ack
4-WAY LBT
WAVELAN CSMA/CA
10 dB Fading Margin
SMALLER NETWORK
INDIVIDUAL STATION PERFORMANCE WITH [redacted] POWER CONTROL

Fig. 11