Power consumption in VEmesh *synchronized flooding* based wireless mesh networks

Executive summary

This paper analyses and compares power consumption between networks based on synchronized flooding to those using routing. Utilizing data from deployed VEmesh systems, it concludes that *synchronized flooding* networks exhibit lower power consumption in large networks or with short messages, such as those used in smart metering and smart lighting; routing based networks are presumably preferable when the number of active nodes is smaller – such as in home appliances networks.

The *synchronized flooding* advantage further improves with the increase in network size - nodes and hops, adding to other distinct advantages of these networks, which include greater robustness, simpler deployment and maintenance and lower cost of ownership.

The paper provides formulas and numeric examples that can assist in choosing the proper wireless mesh network technology for a specific system design.

Contents

Introduction

Wireless mesh networks are an important building block in the emerging "smart-" environments, which include numerous remote devices such as sensors, lamps, actuators, and meters. Within these environments, there is significant demand for solutions with low power consumption, especially in battery powered applications.
In this whitepaper we investigate the topic of power consumption in wireless mesh networks and compare the consumption behaviour of the two main types of networks in use today - *routing* based and non-routing based, which in this paper will focus on *flooding*. We will explain and calculate the power consumption in these two paradigms to help the reader make an educated choice for his or her network use.

Synchronized flooding networks – brief description

In *classic flooding* networks (also called *naïve* flooding), each message from the source node is blindly relayed to all other nodes in the network. This being a mesh network, each node also serves as a repeater, and therefore any node that receives a message will retransmit it, and so on. Each retransmission is called a hop – hence the classification of *multihop* networks. Thus the message advances and spreads through the network in hops, until it is received by all the network's nodes, and eventually by the destined node.

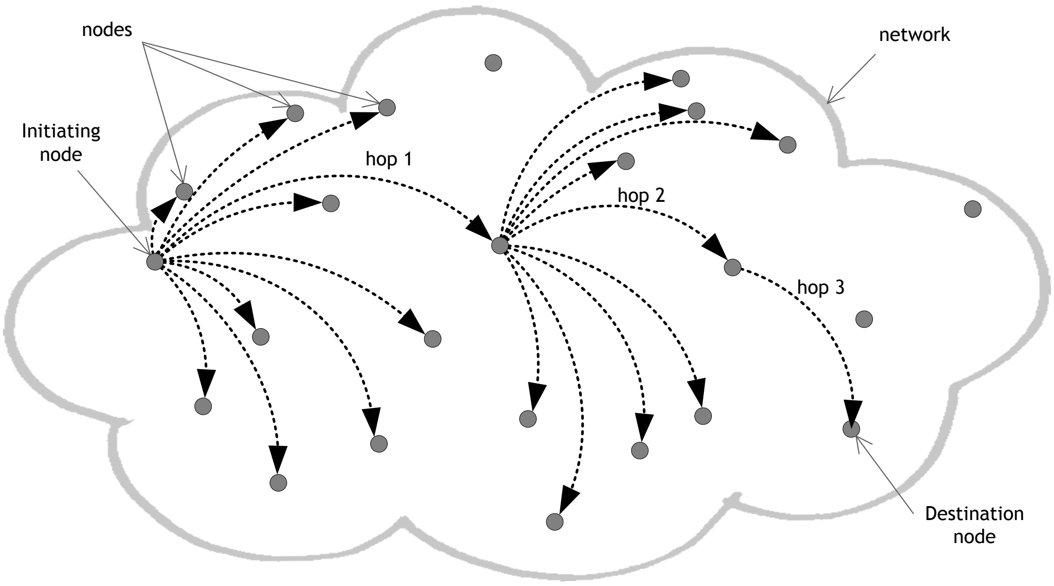

Figure 1 - In a mesh network, the messages propagate in hops (for the sake of clarity, only a few propagations paths are shown)

Since the radio signals propagate in a multitude of paths, and because *flooding* utilizes all available paths, flooding is therefore a highly robust technique to propagate messages through a mesh network. Furthermore, a flooding based network is easy to install and to maintain, with any number of nodes that can be inserted or removed; as long as the added or remaining nodes are within the reception range of at least one other node, the network continues to operate without interruption. Thus, there is no need for self-healing and realignment time, and there is no down-time of the network. Changing the location of a node or of the gateway is also immaterial, as long as the added or remaining nodes are within the reception range of at least one other node.

However, classic flooding networks are also known to suffer from the consequences of several problems, the most critical being the collision and contention of packet messages, which are often referred to as the broadcast storm problem. This problem not only impacts the nodes' power consumption, but also downgrades the latency time and consistency of the network.

A recent paradigm, named *synchronized flooding*, uses an improved approach: all the nodes in the reception range *simultaneously* retransmit the message received from the initiating node - these retransmissions being precisely *synchronized in time*. This is possible in certain modulation schemes like FSK and ASK. The retransmission process is repeated until all nodes (including the farthest node of the network) receive the message, at which time the process ceases.

In *synchronized flooding*, the timing of each hop is precisely defined in a *time slot* of its own, while the time of all the hops required for a message to reach all the nodes is termed *time frame*.

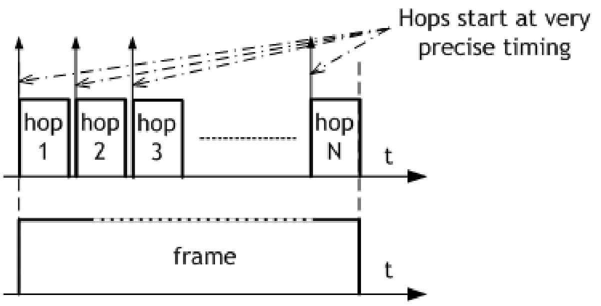

Figure 2 - The precise timing of the hops and the frame

Retransmitting the messages synchronously brings several important advantages, the main one being the solution to the broadcast storm problem. Furthermore, the radio signals that overlap during the reception of a message sum up rather than adversely affect the reception, thus increasing the received energy and hence the overall range. Another advantage regards the network's latency, which becomes small and consistent.

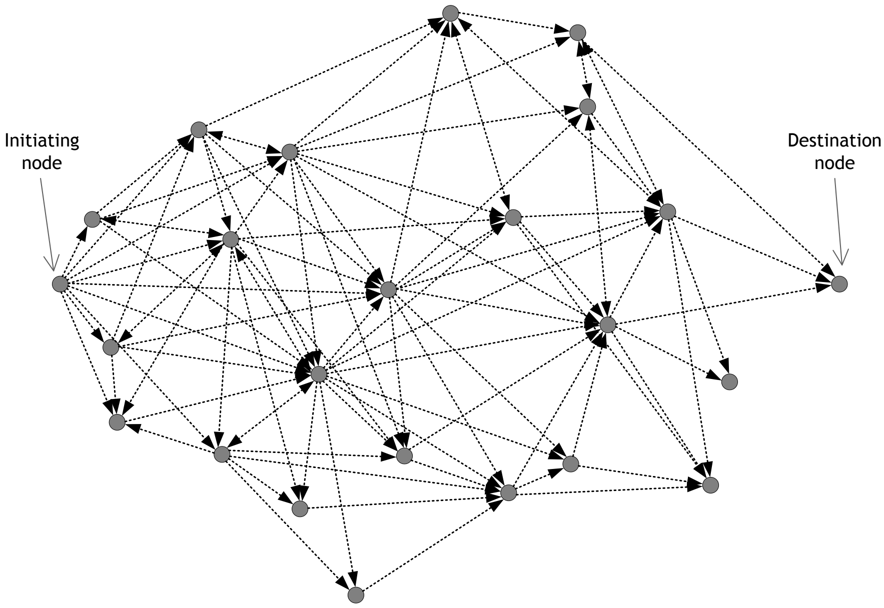

Figure 3 - Illustration of a 3-hop flooding based networks example

Because of its advantages, the synchronized flooding technique is used in all of Virtual Extension's VEmesh networks; its details are beyond the scope of this paper and are described in various Virtual Extension literatures, as well as in several academic papers.

Power consumption factors and computation

Number of participating nodes and average power consumption

This chapter compares the average node power consumption of the two types of networks – *routing* based and *synchronized flooding* based.

Initial intuition suggests that nodes in routing based networks consume, on the average, less power because only a small number of nodes participate in the delivery process of a message; hence the nodes are likely to be, on average, active for *less time*.

However, the average power consumption in the nodes depends not only on how many times they are delivering messages, but also on the *actual power consumption* during their active state, including all internal network processes time that they are active.

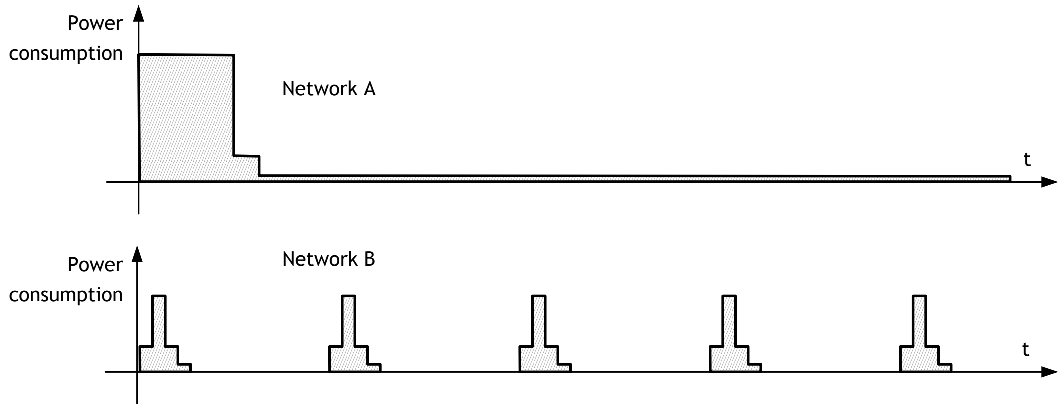

Figure 4 - Illustration of different power consumption patterns for Networks A and B

Assuming that Network A in Figure 4 illustrates the node power consumption in a routing based network and Network B in a synchronized flooding network, we can now intuitively see that their power consumption (represented by sum of the areas marked with diagonal lines) is *comparable*.

For the comparison between the average power consumption of a node, we'll assume that the node acts as a repeater. The reason is that in a typical industrial application using a mesh multihop network, most of the times a node functions as a repeater - rather than as an end node. Furthermore, when acting as a repeater, from the standpoint of power consumption, the node is in its worst case, as it must both *receive* and *transmit*.

Types of nodes and the size of network

One important factor is related to the types of nodes in the network.

A typical *routing* based network has 3 types of nodes: coordinators, routers, and end devices. Amongst them, only the end devices' nodes are designed for low power consumption and can practically be battery operated; for example, the radio receivers on coordinator and routers must be on at all times. On the other hand, the end devices' nodes are not capable of functioning as repeaters - a task reserved for the router nodes.

Hence, the more hops the network has, the more routers are required to participate - compared to the number of end devices' nodes, and as a result the less efficient is a *routing* based network is in terms of power consumption.

A typical *synchronized flooding* network has 2 types of nodes: gateways and regular nodes. Both types are similar in their design and can be designed for low power consumption, as both can operate on low duty cycle (thus saving on reception power) and do not require the processing power and memory of routing processing (thus saving on processing power). Furthermore, all regular nodes can operate as end devices or as repeaters. As a result, in a typical battery operated *synchronized flooding* network, all nodes operate in a duty cycle mode that can be made very low, directly because the transmission times are very short (no message overhead).

In this respect, the *synchronized flooding* advantage is evident in networks that have a larger number of hops, and even more in those that can also have a lower duty cycle.

Node power consumption - transmission

In all mesh networks, each message carries two types of data: (1) the actual payload – the string of actual data to be conveyed between the nodes, and (2) the overhead of network management information. This overhead can range from a few bits or bytes of data per message, to messages whose sole purpose is to pass routing information throughout the network.

When transmitting a message, the node consumes (a) transmitter power and (b) processor and memory power. The total transmission power is closely linked (usually proportional) to the transmission time.

For analysing the transmission time, we'll first define the efficiency of the transmission in relation to the payload. The efficiency is the ratio between the amount of data bits in the payload and the total amount of bits in the message:

$$\varepsilon = \frac{\alpha}{\alpha + o}$$

Where ε is the Efficiency

α is the Payload (in bytes or bits)

o is the Overhead (in bytes or bits – same units as the Payload)

In *routing* based networks, the efficiency has a strong dependence on the 2 types of overhead: (1) management messages, such as those requested for propagating routing information[1], and (2) data in each message.

The overhead size depends on several factors, such as the routing protocol in use, the specific application, the various propagation and temporary conditions, as well as on the number of nodes and hops in the specific network. The latter two factors determine the size and the complexity of the routing tables and increases rapidly with the size of the network (nodes and hops).

In *synchronized flooding* networks the message consists of almost pure data, with few bits per message for proper operation of the physical layer, as there is no routing information to propagate. The amount of bits can be made constant per message[2], there are no management messages, and as a result the efficiency is high and deterministic.

[1] or route request, route reply, and route maintenance

[2] in VEmesh the number of bits depends only on the VEmesh model

Node power consumption - reception

While receiving a message, the node consumes (a) receiver power and (b) processor and memory power. We will analyse each power consumption, separately.

Receiver power consumption

In a typical *synchronized flooding* network, the node's receiver is synchronously ON only during the hop times defined by the network. During this time can a node receive network transmissions, as explained and depicted in Figure 2. After receiving a message (from other nodes or from the gateway), the node – while acting as a repeater – retransmits the same message in the next hop. In the next hop (the hop following retransmission), the node will listen for confirmation that its message has been handed off. Retransmission is confirmed when the node hears its message repeated by its peers.

During the other hops, the node's receiver is usually ON only for a brief period of time, during which it either tries to synchronize with the network and can't (since it doesn't receive enough signal), or receives a signal that does not require retransmission. Hence during these hops, the node's receiver and processor are ON for only a fraction of the hop time.

In a typical *routing* network, the receiver is usually ON for much longer periods since there is no single time in which message reception is expected.

Processing power consumption

The tasks performed by a processor of a node within *routing* based networks usually include receiving (parsing) the routing information attached to the message payload, analysing it, calculating a route, and attaching the new routing information to the transmitted message.

For performing these tasks, the node needs to use a processor with enough computing power. Furthermore, it needs to have enough memory for processing the routing matrices, during which the processor must be in an active state.

The demand of computing power and amount of memory increases when the number of the nodes and the number of the hops increases within the network. As a result, very few networks can use affordable processors to handle networks with more than 10-15 hops when utilizing routing schemes for signal propagation.

In contrast, none of the above tasks is required in the nodes of *synchronized flooding* networks, as there simply is no routing. Therefore, synchronized flooding nodes can be active for less time, and are in an idle or equivalently low power state the rest of the time.

And since *synchronized flooding* nodes can utilize simple processors and very little memory, their power consumption during the active state is further reduced.

Total power consumption

The above chapters explain why the power consumption of a node in a *synchronized flooding* network is always lower than in a *routing* network. To summarize the main points:

1. The *transmission* efficiency of flooding nodes is much higher, since routing nodes have to spend transmission time, and its equivalent power, on routing tables and management, resulting in overhead of dedicated messages and data in messages;
2. The *reception* time in flooding nodes is much shorter, since they are switched ON only for brief periods of time;

3. The *processing* requirements of flooding nodes are much more modest, so the processors and memory are much simpler and hence less power consuming. Furthermore, the flooding nodes are active significantly less time than routing nodes.

Two of these points - *transmission* and *processing*, are related to the routing tables' size. The larger the routing network (nodes and hops), the bigger the routing tables – a significant disadvantage of routing networks' total power consumption.

The next chapter exemplifies the power consumption in a typical wireless mesh network using synchronized flooding.

Calculating the power consumption in synchronized flooding

This chapter analyses the average power consumption of a node Q_H , which functions as a repeater in a network of N_H hops. Note that at the same time, additional nodes might receive the same message, as illustrated in Figure 3. The same analysis applies to them as well.

A good first approximation of the average power consumption of a node in any mesh network is during its activity as a repeater – when it both receives and transmits. In a typical mesh network, the number of hops is $N_H \geq 10$ (see note [3] below); hence, the number of times that a node functions as a repeater is at least one order of magnitude greater than the number of times it functions in any other mode, hence the validity of the approximation.

The power consumption of the node Q_H , while functioning as a repeater, is illustrated in Figure 5 below. As mentioned before, this is a synchronized operation scheme, where each *hop* occurs in its allocated *time slot*, as illustrated in Figure 2 above. The two terms are therefore closely related.

[3] The number of hops N_H in typical mesh applications, such as smart metering, smart illumination, building automation or agriculture is usually between 10 to 40; in Virtual Extension networks the default number of hops is 20.

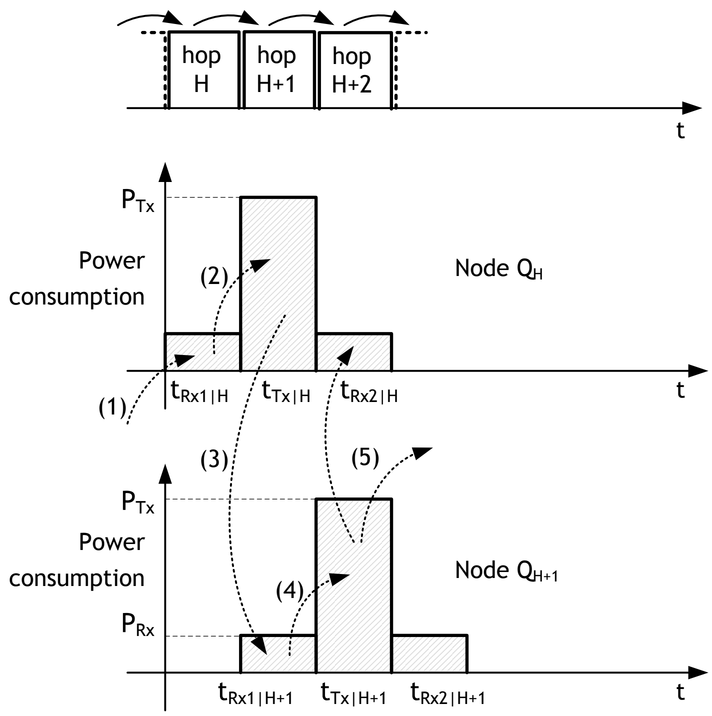

Figure 5 - Activity and power consumption timing - simplified diagram

The time diagram begins when node Q_H enters receive mode and starts to receive a message packet at hop number H, for the first time after the beginning of a new frame. This reception, marked as step (1) in Figure 5 above, occurs during time slot t_{Rx1} of the node Q_H. Other nodes might have their own t_{Rx1}, so for the sake of clarity let's mark the time slot t_{Rx1} of the node Q_H as $t_{Rx1|H}$. While in receive mode ($t_{Rx1|H}$), node Q_H receives the message, processes it, and retransmits the message (following our assumption that the node operates as a repeater). This step is marked in Figure 5 as step (2),

The power consumption of node Q_H during $t_{Rx1|H}$ time slot (hop H) is:

$$P_{Rx} = P_R + P_P \quad (1)$$

Where P_{Rx} is the *total* power consumption of the node during reception,

P_R is the power consumption of the *receiver*,

P_P is the power consumption of the *processor in active mode*.

Functioning as a repeater, the node Q_H retransmits the message during the next time slot, $t_{Tx|H}$, which occurs at hop $H+1$.

During the transmission, the power consumption is:

$$P_{Tx} = P_T + P_P \quad (2)$$

Where P_{Tx} is the *total* power consumption of the node during transmission,

P_T is the power consumption of the *transmitter*,

P_P is the power consumption of the *processor in active mode*.

The transmission of node Q_H (at $t_{Tx1|H}$) is received by at least one node (node Q_{H+1} in this example), and it occurs during the same time - at hop $H+1$, as depicted in step (3) of Figure 5 above. Upon switching to receive mode, the receiver of node Q_{H+1} receives the transmission with message, processes it, and decides to retransmit it, identical to the operation of node Q_H one time slot earlier. The transmission of Q_{H+1} will occur at hop $H+2$, and this step is marked in Figure 5 as step (4).

During hop $H+2$, the node Q_H receives the same message packet it originally transmitted in $H+1$, except that this time the transmission came from Q_{H+1} (and possibly from additional nodes that received its transmission t_{Tx}, at hop $H+1$); this step is marked in Figure 5 as step (5). The reception of this re-transmission provides confirmation that at least one node (in this case Q_{H+1}) had received transmission t_{Tx} by node Q_H, at hop $H+1$.

The power consumption of the node Q_H during time slot $t_{Rx2|H}$ is identical to its power consumption during time slot $t_{Rx1|H}$ ($P_{Rx} = P_R + P_P$).

After receiving the transmission of node Q_{H+1}, the node Q_H enters a passive mode. During this passive mode, the receiver turns ON at the beginning of each time slot, checking for a transmission addressed to it. If no relevant transmission exists, the node goes idle for the remainder of the time slot. When idle, the receiver is off (as is the transmitter), leaving only the processor, which consumes idle power P_I. The passive mode lasts $(N_H - 3)$ hops, after which the frame starts again, as illustrated in Figure 6 below.

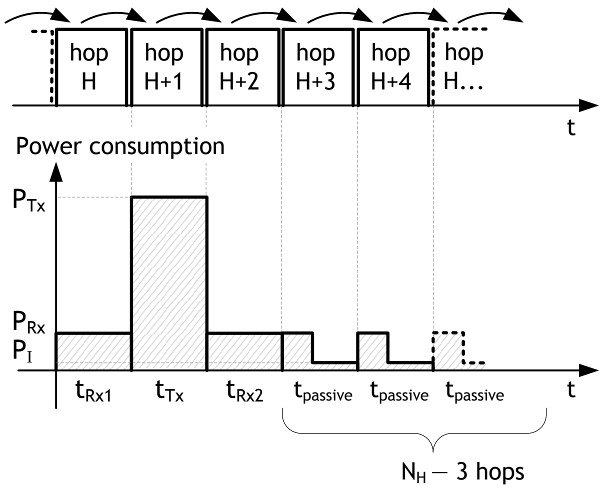

Figure 6 - The passive mode lasts N_H -3 hops and then the cycle starts again

The *average* power consumption of the node in a network with N_H hops is:

$$P_{A(N)} = \frac{P_{Tx} + 2 \times P_{Rx} + (N_H - 3) \times (\eta \times P_{Rx} + (1-\eta) \times P_I)}{N_H} \quad (3)$$

Where η is the alert ratio, defined as the part of the time slot during which the receiver and processor are ON; it can also be described as the internal duty cycle of a $t_{passive}$ time slot. The time $t_{passive} \times \eta$ is necessary for the node to receive and detect a training sequence (if necessary), in order to decide if there is a valid transmission going on.

Correspondingly, $(1-\eta)$ is the remaining part of the time slot during which the receiver is OFF and the processor is in idle mode. The power consumption of the processor during $(1-\eta)$ is P_I ; during this time interval only the processor is ON – in idle mode.

The formula (3) above assumes that a new frame starts again immediately after the end of the N_H hops, thus allowing for the fastest possible response time of the network. In battery operated networks, the new frame never starts immediately after the last. In these types of networks, low power consumption is more important that fast response times, and therefore the networks are set to operate at slower response times in order to reduce power consumption, resulting in an increased battery life. This is a reasonable trade-off, achievable by operating the network in a low duty cycle mode.

In synchronized flooding network, this is attained by delaying the start of the frame after the end of the N_H hops.

A convenient way of expressing the delay is in number of hops (of delay), thus:

$$N_T = N_H + N_D \tag{4}$$

Where N_T is the *total* number of hops,

N_D is the delay, expressed in number of hops, and

N_H is the number of hops in the network.

(The duty cycle of the network is then: $\tau = \dfrac{N_H}{N_T}$).

We can then calculate the total power consumption of the node, to be:

$$P_{A\tau} = \frac{P_{Tx} + P_{Rx} \times (2 + \eta \times (N_H - 3)) + P_I \times (N_D + (1-\eta) \times (N_H - 3))}{N_D + N_H} \tag{5}$$

Where $P_{A\tau}$ is the average power consumption of a node in a network operating at a duty cycle of τ .

As expected, $P_{A\tau}$ decreases asymptotically to the idle power P_I when the delay N_D increases.

Since the voltage is constant, the average current consumption can be expressed the same way:

$$I_{A\tau} = \frac{I_{Tx} + I_{Rx} \times (2 + \eta \times (N_H - 3)) + I_I \times (N_D + (1-\eta) \times (N_H - 3))}{N_D + N_H} \tag{6}$$

Where: I_{Tx} is the total power consumption (transmitter and processor) during transmission,

I_{Rx} is the total power consumption of the node during reception,

I_I is the power consumption of the processor when idle,

$I_{A\tau}$ is the average power consumption of the node.

Example of power consumption in a synchronized flooding network

This chapter shows a numerical example, based on the simplified calculations above. The example is of a real-life battery operated synchronized flooding network, used in a smart water metering application.

The specifications of the network in this example are as following:

- The maximum payload is 10 Bytes: $P_B = 10B$
- The number of nodes (meters) is: $M = 200$;

- The number of hops is: $N_H = 20$
- The minimum frequency of reading the meters is twice a day: $t_{M(\min)} = 12h$
- The nodes use C-size batteries with a typical capacity of $C_B = 8Ah$

Where the battery capacity is chosen to enable the system to operate for the required time duration:

$$C_B = I_{A\tau} \times L \quad (7)$$

In this example, the requirement is of minimum life time of the batteries is 7 years: $L_{(\min)} = 7Y$

Then the equation (7) can be used to calculate the maximum average current for this example:

$$I_{A\tau(\max)} = \frac{C_B}{L_{(\min)}} = \frac{8Ah}{7Y} \cong 0.13mA$$

This example makes use of a typical VEmesh[4] synchronized flooding network, with the following characteristics:

- The data bitrate is: $B = 50Kbps$.
- The transmitter current consumption is: $I_T = 26mA$ (at 10mW transmission power)
- The processing current consumption is: $I_P = 2.5mA$
- The receiver active current consumption is: $I_R = 3.5mA$
- The receiver idle current consumption is: $I_I = 6\mu A$
- The alert ratio (in the $t_{passive}$ time slot) is: $\eta = 0.3$

Using the values above, we can calculate the minimum time of a hop to be:

$$t_H = \frac{P_B}{B} \quad (8)$$

Which in this example is $t_H = \frac{10 \times 8}{50 \times 10^3} = 1.6ms$

We can then calculate the minimum time of interrogating all the M nodes, when using a simple sequential polling. The time is:

$$t_M = 2 \times t_H \times N_H \times M \quad (9)$$

Which in this example is: $t_M = 2 \times 1.6 \times 10^{-3} \times 20 \times 200 = 12.8s$

These minimum results were obtained while operating at a duty cycle of 100%. This being a battery operated application, it is operated at a low duty cycle as to increase the battery life at the expense of a slower response time - as explained in the previous chapter.

[4] VEmesh are Virtual Extension wireless mesh networks that use synchronized flooding

Using the equations devised above: (6), (7), (8) and (9), and varying the amount of duty cycle N_D, we can plot a graph that represents the dependence between the response time t_M and the battery life L in this example:

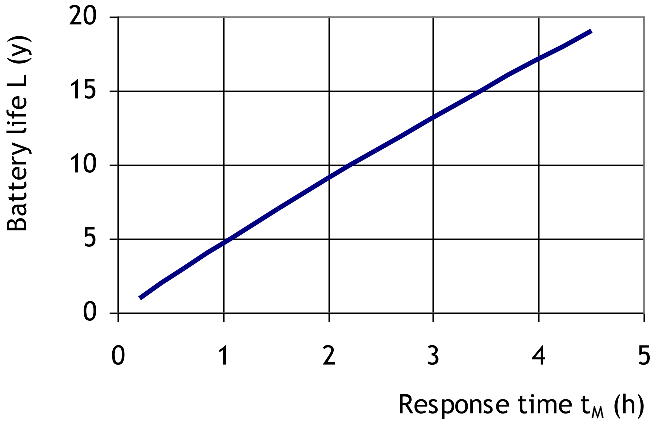

Figure 7 - Battery life versus the response time in the example

The graph in Figure 7 above shows that the requirements of the application are easily met. The life time $L_{(\min)}$ of 7 years can be achieved by receiving the reading of all the network nodes as frequently as every 1½ hours ($t_M = 1.5h$) - much faster than the original requirement of every 12 hours $t_{M(\min)} = 12h$, even when using a simple polling technique[5]. Alternatively, the life time can be increased.

Comparison between power consumptions

We will now compare data from the above computations with data from one of the very few papers containing real power consumption numbers in a routing based network[6] (the scarcity of this data, mentioned by the author and according to him, is because "*most of such* [power consumption] *related articles are proprietary for the ZigBee Alliance members and not otherwise available"*).

The data presented in the routing paper was compiled by the author in the laboratory and validated in an industrial manufacturing environment, using off-the-shelf commercial components[7].

The measurements' set in the routing paper uses a time interval of 25,000 msec. $(t_H \times N_H)$, for 2 cases: (1) when the number of hops N_H is 1 and (2) when the number of hops N_H is 2. For comparison purposes, only the 2-hops case is used, as it is more characteristic of real-life mesh networks[8].

[5] even better results are obtained by using more elaborated techniques, but their description is beyond the scope of this paper

[6] Chaitanya S. Misal "Analysis of power consumption of an end device in a ZigBee mesh network" thesis, the faculty of The University of North Carolina at Charlotte, 2007.

[7] Maxstream™ ZigBee™ Series 2 modules (Maxstream firmware)

[8] The number of hops N_H in typical mesh applications, such as smart metering, smart illumination, building automation or agriculture is usually between 10 to 40; in Virtual Extension networks the default number of hops is 20.

For the 2-hops case, the routing paper measurements were performed using several sub-cases, split according to the number of bytes in the package. Two of them are relevant: (1) when the number of bytes P_B is 24[9] and (2) when the number of bytes P_B is 48[10].

Finally, the electronics in the routing paper uses roughly the same voltage as VEmesh components (3V).

The resulting current consumption from both systems – VEmesh and the routing system, using the above data set, are consolidated in the following table:

P_B	VEmesh	Routing	Ratio
24B	22.6 µA	221.6 µA	4.9
48B	33.2 µA	236.8 µA	3.6

Table 1 - Comparison between the current consumption of VEmesh versus routing based

Table 1 above shows that the power consumption of a general purpose VEmesh network is *significantly lower* than of an similar routing network, under the conditions described in the above routing paper: 4.9 times better in one case and 3.6 times better in the other case.

We'll summarize this chapter noting that its results validate the previous conclusions regarding the power consumption advantage of *synchronized flooding* approach compared to routing, as well as the conclusion that the lower the number of bytes in the message, the greater the advantage. Based on the previous chapters, we can safely assume that in a larger network (with more hops and nodes), the advantage of the *synchronized flooding* approach is even more significant.

It should also be noted that the routing system for the comparison is a ZigBee™ based system, and in the market exist other *proprietary* routing based system which likely have better performance regarding power consumption. On the other hand, the VEmesh™ system in the comparison is also general purpose, and can be further tuned and optimized for power consumption performance, thus leaving the balance in favour of VEmesh.

Conclusions

This paper explains, exemplifies, and demonstrates, using a real-life example, that the power consumption of *synchronized flooding* networks has a clear advantage regarding power consumption performance in typical or large sized networks. In small networks, routing based networks might consume less power. The break-even point depends on the network characteristics, with the simplified numerical example above serving as a practical indication.

Diversity Path Mesh™ is Virtual Extension's implementation of *synchronized flooding* in its VEmesh wireless mesh networks. Aware of the significance of low power consumption in its mesh networks and committed to continuously improve the already high quality of its products, Virtual Extension constantly adds new and improved ways of increasing battery life, both by using new and improved algorithms, as well as diversifying the components in use.

[9] Table 5-9 on page 70 in Chaitanya S. Misal paper
[10] Table 5-10 on page 72 in Chaitanya S. Misal paper

Virtual Extension, 5 Menachem Begin Blvd., Beit Dagan 50200, Israel +972-3-7321207
Website: www.virtual-extension.com Email: info@virtual-extension.com