Optimized Design of a Human Intranet Network

Ali Moin[1] Pierluigi Nuzzo[2] Alberto L. Sangiovanni-Vincentelli[1] Jan M. Rabaey[1]
[1] Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
[2] Department of Electrical Engineering, University of Southern California, Los Angeles
Email: {moin, alberto, jan}@eecs.berkeley.edu, nuzzo@usc.edu

ABSTRACT

We address the design space exploration of wireless body area networks for wearable and implantable technologies, a task that is increasingly challenging as the number and variety of devices per person grow. Our method efficiently decomposes the problem into smaller subproblems by coordinating specialized analysis and optimization techniques. We leverage mixed integer linear programming to generate candidate network configurations based on coarse energy estimations. Accurate discrete-event simulation is used to check the feasibility of the proposed configurations under reliability constraints and guide the search to achieve fast convergence. Numerical results show that our application-specific approach substantially reduces the exploration time with respect to generic optimization techniques and helps provide clear identification of promising solutions.

1. INTRODUCTION

Wearable and implantable devices are destined to become key components of the Internet of Things (IoT), with a myriad of applications, ranging from health-care and sport to entertainment and communication. Networks of these devices, often called Wireless Body Area Networks (WBANs), have been extensively investigated in the literature [1, 2]. However, most of the previous work on WBANs focused on a limited number of nodes, often targeting a specific application. With the explosive growth of the wearables market, having tens of these interconnected devices per person, performing different functions, is not unreasonable. Further, a general approach to the design of a WBAN architecture with multiple, heterogeneous nodes is needed to shorten design and deployment time. We call this architecture *Human Intranet* [3]. A Human Intranet should seamlessly integrate an ever-increasing number of sensor, actuation, computation, storage, communication, and energy nodes located on, in, or around the human body, and acting in symbiosis with the functions provided by the body itself.

As the number and variety of nodes, network protocols, and technologies available at each layer continuously increase, the design of such an open, adaptive, and scalable network platform becomes a daunting task. In a Human Intranet, design requirements often span the entire envelope of achievable performance. When a safety-critical node such as a wearable insulin delivery device is part of the network, reliability becomes of utmost importance. Conversely, for an everyday physical activity monitoring application, achieving the longest possible battery lifetime is preferred, while a few packet drops can occasionally be tolerated. Since key design concerns such as reliability and lifetime may be conflicting and depend in an inextricable way on the architecture and layers of the network, design aids that help explore a large design space and customize the network for different applications are sorely needed.

This paper proposes a methodology for the design space exploration (DSE) of a Human Intranet. Our contributions can be summarized as follows:

- A formulation of the DSE problem in terms of optimized mapping of system requirements into an aggregation of components from a library that encompasses all the network layers. Our goal is to select network configurations that maximize lifetime under constraints on reliability. To the best of our knowledge, this is the first DSE framework for WBANs enabling the exploration of configurations across the entire network stack.
- An efficient algorithm that decomposes the DSE problem into smaller subproblems by coordinating specialized analysis and optimization methods. We leverage mixed integer linear programming to generate candidate network configurations based on coarse energy estimations. Accurate, discrete-event simulation is used to check the feasibility of the proposed configurations under reliability constraints, and guide the search to achieve fast convergence. On our problem instances, our algorithm outperforms state-of-the-art general-purpose optimization methods, such as simulated annealing.

We demonstrate the effectiveness of our formulation and algorithm with an example. In the example, we highlight the impact of different network topologies on the system performance metrics and the substantial dependency of the selected network configuration on the target application.

Related Work. Human Intranets are special cases of Wireless Sensor Networks. DSE methods for Wireless Sensor Networks (WSNs) can be categorized in two classes, based on the approach used to evaluate the system performance

DAC '17, June 18 - 22, 2017, Austin, TX, USA

ISBN 978-1-4503-4927-7/17/06…$15.00

DOI: http://dx.doi.org/10.1145/3061639.3062296

metrics for different design choices [4]. Simulation-based approaches tend to be more accurate but considerably slower, since they require running lengthy simulations on complex, high-fidelity models. On the other hand, analytical approaches provide faster ways of evaluating the network performance but are less accurate due to the simplifying assumptions usually adopted to obtain tractable models.

Beretta et al. [5] explore the energy-performance tradeoffs of a WBAN using an accurate analytical model that is orders of magnitude faster than simulation. However, their focus is on a specific network configuration, a star topology with Time Division Multiplexing (TDMA) Media Access Control (MAC) layer from the IEEE 802.15.4 standard [6]. We target, instead, a larger design space for which an accurate compact model is not available. We then combine a coarse analytical model, to rapidly prune the design space and accelerate the search, with a smaller number of performance evaluations via simulation, to achieve high accuracy on a reduced set of promising candidates.

In Grassi et. al. [7], the authors encapsulate domain-specific knowledge and rules into a discrete-state Markov decision process that is used to navigate the design space. The method is applied to IEEE 802.15.4 star networks and is shown to reduce substantially the number of simulations needed to converge. Integrating domain knowledge within the decision process is desirable for efficient design space exploration. However, it is not straightforward to implement at the early stages of the design process, especially when the design space spans different topologies and MAC layer configurations, and in the presence of high temporal variations of the WBAN channel. We leverage the approach of Nuzzo et. al. [8, 9] and Finn et. al. [10] characterized by the coordination of Mixed Integer Linear Programming (MILP) for design space optimization with simulation-based branching to incorporate Human Intranet domain-specific knowledge and steer the MILP search toward the most promising solutions. Our approach is general, since it applies to different classes of models and can be integrated with data-driven, automated model construction techniques [4].

2. PROBLEM FORMULATION

We follow a platform-based design methodology [11] to cast the design problem as a mapping problem where system specifications are mapped to an architecture, which consists of the interconnection of components available in a library.

2.1 HI Architecture

As shown in Fig. 1, an HI consists of N heterogeneous sensing, actuation, and processing nodes that share a wireless channel around the body. We assume that nodes may be placed in M pre-determined locations. We use a binary variable n_i to indicate the presence of a node at location i, i.e., $n_i = 1$ if and only if location i is used, with $i \in \{0, \ldots, M-1\}$. We can then represent a network *topology* as a vector $\nu = (n_0, \ldots, n_{M-1})$ specifying the locations of the nodes, under the integer linear constraint $N = \sum_{i=0}^{M-1} n_i$. Mixed integer linear inequalities can also be used to specify additional topological constraints. For instance, we can require that location i be used if location j is used by writing $n_j - n_i \leq 0$.

2.1.1 Wireless Channel

The wireless channel model provides the instantaneous

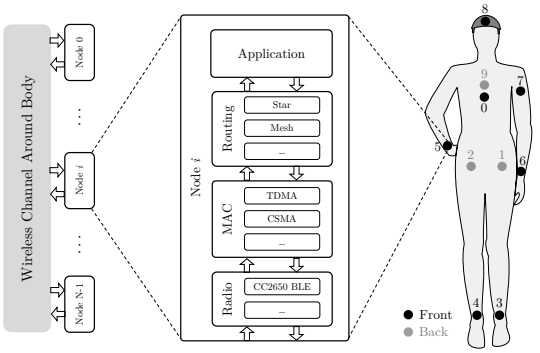

Figure 1: Diagram of a Human Intranet

path loss (in dBm) between any pair of node locations (i, j), that is, the attenuation of the power density of an electromagnetic wave traveling between the two locations. We use a probabilistic model, based on [12, 13], in the following form:

$$PL_{i,j}(t) = \overline{PL}_{i,j} + \delta_{PL_{i,j}}(t), \quad (1)$$

where $\overline{PL}_{i,j}$ is the average path loss over time and $\delta_{PL_{i,j}}(t)$ captures the temporal variation. The average path loss can be inferred from measurements on human subjects for a predefined set of locations on the body. The temporal variation is, instead, a random variable capturing the effect of body movements and environmental changes on the channel quality. $\delta_{PL_{i,j}}(t)$ at time t is drawn out of a probability density function that depends on the observed value $\delta_{PL_{i,j}}(t - \Delta t)$ at time $t - \Delta t$ and the time Δt elapsed since the last observation. Intuitively, if little time has passed, $\delta_{PL_{i,j}}(t)$ does not significantly differ from $\delta_{PL_{i,j}}(t - \Delta t)$. The actual probability density functions can also be empirically estimated from measurement data.

2.1.2 Node

Each node supports the four standard networking layers. The radio (physical layer) is responsible for the transmission and reception of raw bit streams over the wireless medium. The MAC (data link layer) mechanism regulates how nodes gain access to the shared channel. The routing mechanism (network layer) manages addressing and routing between nodes. Finally, the application (application layer) is in charge of resource management and provides the high-level Application Programming Interface (API) to the node.

Radio. The HI radio is characterized by the following configuration parameters

$$\chi_{rd} = (f_c, BR, Tx_{dBm}, Tx_{mW}, Rx_{dBm}, Rx_{mW}), \quad (2)$$

where f_c is the carrier frequency, BR is the bit rate, Tx_{dBm} and Tx_{mW} are, respectively, the transmitter output power and power consumption, Rx_{dBm} and Rx_{mW} are the receiver sensitivity and power consumption.

For successful packet reception, we require that the transmitter output power (in dBm) of a node be at all times larger than the sum of the receiver sensitivity and the path loss between any pair of locations i and j occupied by nodes. Formally, $Tx_{dBm} \geq Rx_{dBm} + PL_{i,j}(t)$, $\forall\ t \geq 0$, $\forall\ (i, j) : n_i = n_j = 1$.

Based on the bit rate BR, we can evaluate the *packet transmission duration* as $T_{pkt} = \frac{8L}{BR}$ (in seconds), where we assume that each physical layer packet contains L bytes.

Finally, the *power consumption* used to transmit and receive packets can be estimated as follows. Most modern radios stay in sleep mode by default, and wake up either periodically or using a wake-up receiver to send or receive packets. We assume a periodic traffic with the same throughput, transmitter, and receiver power consumption for all nodes. At each round of transmissions, each node is then involved in one transmission event and $N-1$ reception events yielding:

$$P_{rd/tx} = Tx_{mW} + (N-1)Rx_{mW}. \quad (3)$$

Media Access Control. The access control mechanism provides the nodes with access to the shared channel while avoiding packet collisions, by using either a contention-based or a time-multiplexed strategy. In our library, we can choose between a Carrier Sense Multiple Access (CSMA) and a Time Division Multiple Access (TDMA) protocol, which are then implemented in the discrete-event network simulator. The MAC configuration vector $\chi_{MAC} = (P_{MAC}, B_{MAC}, AM, T_{slot})$ includes the binary variable P_{MAC}, encoding the selected protocol, the integer variable AM, encoding the CSMA access mode, the buffer size B_{MAC}, and the time-slot duration T_{slot} (in seconds) for the TDMA protocol.

Remark. In CSMA, the transmitting node senses the medium to make sure that no other transmission is in progress and initiates the transmission by following a protocol given by the access mode. This causes a non-deterministic delay in communication. In TDMA, each node has exclusive access to the medium during its dedicated time slot, which makes the communication deterministic. However, scheduling the access for all nodes requires maintaining a global synchronized clock, which can be challenging, especially in an *ad hoc* wireless network with no central coordinator in charge of this task.

Routing Mechanism. We can choose between a *star* and a *mesh* network topology in our library. Therefore, the routing configuration vector $\chi_{rt} = (P_{rt}, n_{coor}, N_{hops})$ includes a binary variable P_{rt} encoding the selected protocol ($P_{rt} = 1$ if *mesh* is selected and 0 otherwise), a variable n_{coor} defining the central coordinator node in a star topology, and the maximum hop count N_{hops} for a mesh topology.

WBANs traditionally use a *star* network topology [1, 2], often with a central hub gathering bio-signals from several sensor nodes around the body, as is also included in the IEEE 802.15.6 WBAN standard [14]. However, while being attractive because of its simplicity and low energy consumption, a star topology might not be the best solution for a highly dynamic network with tight reliability constraints, because of the higher probability for the nodes to miss packets as a result of deep fading. An alternative multi-hop *mesh* topology, establishing redundant parallel links between nodes, may then be preferred. This topology would eliminate the need for a central coordinator hub, and allow for a fully distributed network.

Mesh networks generally relay messages using either flooding or point-to-point forwarding schemes [15]. We opt for a controlled flooding algorithm as a suitable choice for the highly dynamic Human Intranet, which avoids the overhead of constant routing discovery as in point-to-point forwarding schemes. In flooding, each node rebroadcasts any received packet if it is not the packet's final destination. To prevent infinite circulation of duplicate packets, the packet payload contains a hop counter which increments every time a node is visited and blocks further retransmissions after N_{hops} is reached. Moreover, a payload includes a history of the nodes reached by the packet to avoid revisiting the same node. This causes each packet to be transmitted at most a finite number of times, denoted as N_{reTx}.

Application Layer. The application layer abstracts all the sensing, actuation, or processing functions of a node. At this layer, we keep track of the sequence numbers of sent and received packets to monitor and evaluate the network performance and power consumption. We parametrize this layer using the configuration vector $\chi_{app} = (P_{bl}, L_{pkt}, \phi)$, where P_{bl} is the node baseline power consumption due to all the node components other than the radio circuitry, L_{pkt} is the length of the generated packets, and ϕ is the data throughput (in packets per second), which is assumed equal for all nodes.

2.2 HI Performance Metrics

We characterize the performance of the network in terms of lifetime and packet delivery ratio.

Network Lifetime. The network lifetime is the time taken for the first node in the network to run out of energy. This is very critical to the autonomy of a BAN, where we aim to minimize the frequency of changing or recharging the battery of the wearable or implantable nodes, or maximize the effectiveness of energy harvesting.

If P_i is the power consumption of node i, and $E_{bat,i}$ is the total stored energy in that node, the network lifetime NLT is defined as

$$NLT = \min_{0 \le i \le N-1} \left\{ \frac{E_{bat,i}}{P_i} \right\}. \quad (4)$$

P_i can be computed as the sum of the baseline power P_{bl} and the radio power consumption P_{rd} for all nodes. While our design procedure relies on accurate computation of the overall power consumption via a discrete event network simulation, we also provide an approximate model for P_{rd} which we will use to generate promising candidate configurations to be simulated. Using (3), we obtain:

$$P_{rd} = \begin{cases} \phi T_{pkt}(Tx_{mW} + 2(N-1)Rx_{mW}) & P_{rt} = 0 \\ \phi T_{pkt} N_{reTx}(Tx_{mW} + (N-1)Rx_{mW}) & P_{rt} = 1 \end{cases} \quad (5)$$

where we consider ϕ packet transmissions per second, each lasting T_{pkt} seconds. For each transmission event, in a star topology ($P_{rt} = 0$), nodes can both receive the original packet and its retransmitted copy by the coordinator, which motivates the factor of 2 in (5). In a mesh topology ($P_{rt} = 1$), retransmissions are captured by the factor N_{reTx}.

Packet Delivery Ratio. The Packet Delivery Ratio (PDR) is the probability that a generated packet reaches its final destination, which is an indicator of the reliability level of the network. We use a discrete event network simulator to estimate the PDR value, while including stochastic and second-order effects that cannot be captured by a compact analytical model.

To do so, we exploit the packet sequence numbers used by the application layer to track the statistics of the packets sent and received for each node in the network. Assume that $N^{(s)}_{i \to k}$ is the number of unique packets sent from node i to node k, during a simulation run, without counting retransmissions, and $N^{(r)}_{i \to k}$ is the number of unique packets received by node k from node i. Then the PDR of node k

can be estimated as:

$$PDR_k = \frac{1}{N-1} \cdot \sum_{\substack{i=0 \\ i \neq k}}^{N-1} \frac{N_{i \to k}^{(r)}}{N_{i \to k}^{(s)}}, \tag{6}$$

where the duration of a simulation run T_{sim} is selected to guarantee that the error between (6) and the desired probability is bounded by a positive tolerance ϵ. We can then define the overall network PDR as the average of the node PDRs:

$$PDR = \frac{1}{N} \cdot \sum_{j=0}^{N-1} PDR_j. \tag{7}$$

2.3 The Optimal Mapping Problem

Network lifetime and reliability requirements tend to conflict when higher transmission powers or redundant transmissions are used to achieve high reliability, since these result in larger power consumption. Motivated by the exploration of these trade-offs, we formulate the design problem as an optimization problem, where we search for a node configuration ν and parameter vector $\chi = (\chi_{rd}, \chi_{MAC}, \chi_{rt}, \chi_{app})$ that maximize the network lifetime subject to a lower bound PDR_{min} on the PDR. Our network is also subject to a set of *topological constraints*, specifying feasible network configurations, and *configuration constraints*, specifying feasible node configurations, e.g., lower and upper bounds on the configuration parameters for different topologies. Both the topological and configuration constraints can be lumped into the component-wise vector inequalities $r_{\mathcal{T}}(\nu, \chi) \leq 0$ and $r_{\chi}(\nu, \chi) \leq 0$, respectively. Formally,

$$\mathcal{P} := \max_{\nu, \chi} \quad NLT(\nu, \chi) \tag{8a}$$

$$\text{s.t.} \quad r_{\mathcal{T}}(\nu, \chi) \leq 0, \tag{8b}$$

$$r_{\chi}(\nu, \chi) \leq 0, \tag{8c}$$

$$PDR(\nu, \chi) \geq PDR_{min} \tag{8d}$$

Based on the expressions in Sec. 2.1 and 2.2, constraints (8b) and (8c) can be posed in terms of mixed integer linear constraints. On the other hand, the objective function in (8a) and the PDR in (8d) can only be evaluated from expensive simulations. Consequently, we propose an algorithm that exploits the structure of Problem (8) to decompose its solution into smaller tasks that can cooperate to rapidly navigate the design space and substantially decrease the number of simulations needed to achieve the optimum.

3. PROPOSED ALGORITHM

We tackle Problem (8) by using an iterative approach that coordinates a MILP solver and a discrete-event network simulator, as illustrated in Figure 2. At each iteration, the MILP solver suggests a set of candidate solutions that satisfy all the topological and configuration constraints while maximizing an approximate expression of the NLT. A set of simulations are then run for those candidates to calculate an accurate value for the PDR, based on (6) and (7), and the objective function. If the reliability constraint is not satisfied, i.e., $PDR < PDR_{min}$, the current candidate configurations are rejected, and the MILP solver is queried to provide a new set of possible solutions at the next iteration. Our intent is to combine the efficiency of state-of-the-art MILP algorithms with the accuracy of simulation.

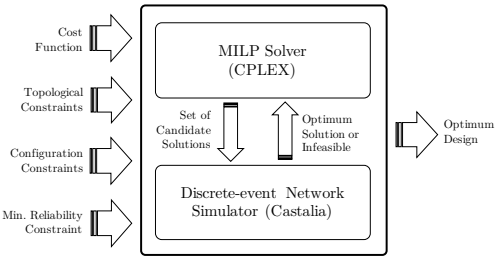

Figure 2: Proposed simulation-based optimization scheme.

Algorithm 1 Human Intranet Design Space Exploration

Input: $\mathcal{P}$
Output: $\mathcal{S}^* = (\nu^*, \chi^*)$

1: $i \leftarrow 0$, $\bar{P}_{min} \leftarrow \infty$
2: **while** true **do**
3: $(\mathcal{S}, \bar{P}^*) \leftarrow$ RunMILP($\tilde{\mathcal{P}}$)
4: **if** $\mathcal{S} = \{\}$ and $\bar{P}_{min} = \infty$ **then return** `infeasible`
5: **else if** $\frac{\bar{P}^*}{\alpha(\mathcal{S}^*, PDR_{min})} > \bar{P}_{min}$ or $\mathcal{S} = \{\}$ **then**
6: **return** $\mathcal{S}^*$
7: $(PDR_{sim}, \bar{P}_{sim}) \leftarrow$ RunSim($\mathcal{S}$)
8: $(status, \mathcal{S}^{**}, \bar{P}^{**}) \leftarrow$ Sort($PDR_{sim}, \bar{P}_{sim}$)
9: **if** $status =$ `feasible` and $\bar{P}_{min} \geq \bar{P}^{**}$ **then**
10: $\bar{P}_{min} \leftarrow \bar{P}^{**}$; $\mathcal{S}^* \leftarrow \mathcal{S}^{**}$
11: $\tilde{\mathcal{P}} \leftarrow$ Update($\tilde{\mathcal{P}}, \bar{P} > \bar{P}^*$); $i \leftarrow i + 1$;

To do so, a key step is to derive a mixed integer linear approximation of the cost function in (8) that can be used by the MILP solver. We first observe that the minimum in (4) is achieved by a non-coordinator node, since a coordinator node (in a star topology) is usually equipped with more energy storage to perform its function. We further assume that all other nodes have the same power consumption $\bar{P}$ and stored energy $\bar{E}_{bat}$, which is a constant. Then, maximizing the network lifetime is equivalent to minimizing the power consumption $\bar{P}$ of a single non-coordinator node. Finally, a mixed integer linear expression for $\bar{P}$ can be found based on (5):

$$\begin{aligned} \bar{P} = P_{bl} + \phi T_{pkt}[(1 - P_{rt})(Tx_{mW} + 2(N-1)Rx_{mW}) \\ + P_{rt} N_{reTx}(Tx_{mW} + (N-1)Rx_{mW})]. \end{aligned} \tag{9}$$

We use (9) to create a relaxed version of Problem (8), which we call $\tilde{\mathcal{P}}$, including all the constraints except for (8d). $\tilde{\mathcal{P}}$ is a mixed integer linear program. We can then combine the solution of this subproblem with accurate simulation-based analysis of the performance metrics to solve the original problem $\mathcal{P}$. Our method is summarized in Algorithm 1.

At iteration i, RunMILP solves the relaxed MILP problem and returns a set of optimal solutions $\mathcal{S} = \{(\nu_j^*, \chi_j^*), 0 \leq j < S\}$ with minimum power consumption according to (9) (line 3 of Algorithm 1). S can be greater than zero, since multiple configurations can minimize (9).

A discrete-event network simulator with an accurate model of the network layers, including the time-varying probabilistic channel path loss, complements the MILP solver. RunSim takes as input the set of candidate solutions $\mathcal{S}$ and returns the vectors PDR_{sim} and $\bar{P}_{sim}$ with simulated PDR and power consumption values for all the configurations in $\mathcal{S}$ (line 7). Sort orders the feasible solutions, that is, the solutions that satisfy the reliability constraint, ac-

cording to their power consumption (line 8). If no solution meets the desired reliability level, a new constraint on the lower bound of $\bar{P}$ is added (line 11) to prune the current set of solutions away from the search space.

If a feasible configuration with minimum power is found then the current solution and cost value $\bar{P}_{min}$ are updated. The algorithm terminates when either the MILP becomes infeasible, i.e., there are no further candidate solutions satisfying the topological and configuration constraints, or when the simulated power consumption of any candidate solution given by the MILP is guaranteed to be higher than the current minimum power value $\bar{P}_{min}$ (line 5). In both cases, the current optimal configuration $\mathcal{S}^*$ is returned.

Since $\bar{P}^*$ provides an estimation of the power consumption under the simplifying assumption that all the messages are correctly received and the retransmissions are successfully performed, it does not represent the lowest possible power consumption that can be obtained by simulating the network for a given PDR_{min}. Therefore, to guarantee optimality, we divide $\bar{P}^*$ by α to formulate the termination criterion in line 5. For a given PDR_{min}, α accounts for any reduction in the simulated power consumption due to the loss of packets. α can be computed as the ratio $\frac{\bar{P}}{\bar{P}_{lb}}$, where $\bar{P}_{lb}$ is the minimum power that a node must consume for the specified PDR bound. If $\frac{\bar{P}^*}{\alpha(\mathcal{S}^*, PDR_{min})} > \bar{P}_{min}$ then no further simulations need to be run since we are guaranteed that new simulations cannot achieve better power consumptions than the current best solution.

Overall, while candidate configurations are generated at each iteration based on an approximate cost function, the selection of the optimum configuration is always based on accurate evaluations obtained by simulation. The algorithm guarantees optimality, since it terminates when no more candidate configurations are available, or additional simulations cannot provide lower power consumption than the current optimum solution.

4. DESIGN EXAMPLE

We demonstrate the performance of the proposed algorithm on a Human Intranet scenario containing various wearable medical devices designed for monitoring human vital signs with applications in both fitness monitoring and medical diagnostics [16]. The MILP routine uses the CPLEX [17] solver via the PYTHON interface PULP [18]. CASTALIA [19], based on the OMNET++ open-source network simulator library, is used as the discrete-event simulator. The implementation of our optimization framework is available online[1]. All the optimization runs were executed on a 2.4-GHz Intel Core i7 with 8 GB of RAM. We set the duration of each simulation to $T_{sim} = 600$ s and averaged the performance metrics over 3 runs to mitigate the effect of randomness. These settings were sufficient to obtain performance estimates within 0.5% relative error.

4.1 Experiment Formulation

We consider the scenario in Figure 1, including 10 potential node positions ($M = 10$), located at the chest, left and right hip, left and right ankle, left and right wrist, left upper arm, head, and back. $\overline{PL}_{i,j}$ for each pair of nodes is based on a two-hour measurement data set capturing the daily activity of adult subjects [20]. The probability density function of $\delta_{PL_{i,j}}$ is adapted from the set of empirical data in [19].

Table 1: TI CC2650 radio specifications

$\mathbf{f_c}$	2.4GHz	**Tx Mode**	$\mathbf{Tx_{dBm}}$	$\mathbf{Tx_{mW}}$
BR	1024kbps	${p_1}^2$	-20	9.55
$\mathbf{Rx_{dBm}}$	-97	${p_2}^2$	-10	11.56
$\mathbf{Rx_{mW}}$	17.7	p_3	0	18.3

A set of topological constraints can be directly formulated based on the application requirements. For instance, one node must be placed on the chest ($n_0 = 1$) for respiration rate monitoring as well as the coordination (n_{coor}) in a *star* topology. At least one node should be at the hip ($n_1 + n_2 \geq 1$) and one at the foot ($n_3 + n_4 \geq 1$) for gait analysis. Finally, at least one node should be placed at the wrist ($n_5 + n_6 \geq 1$) to gather several biological signals, including temperature, heart rate, pulse oxygenation, and motion signals. In addition to these four nodes, we allow up to two more nodes to be arbitrarily placed to possibly improve the mesh connectivity.

A second set of network layer configuration constraints originate from the available protocols and configuration settings, as discussed in Section 2.1. We use the CC2650 radio chip from Texas Instruments [21], a state-of-the-art radio chip for WSN applications, whose parameters are summarized in Table 1. For example, we encode the selection of one of the three transmission power levels using the binary variables p_1, p_2, and p_3. We then add the constraint $p_1 + p_2 + p_3 = 1$ to state that only one power value is possible in each configuration.

The MAC layer implementation in the discrete event simulator supports both the CSMA and TDMA options. For CSMA we use a *TunableMAC* implementation with non-persistence access mode [19], which reduces the collisions by backing off for a random amount of time if the medium is busy. For TDMA, we use 1 ms time slots assigned equally to all nodes in round-robin fashion. The routing constraints can also be captured by a set of mixed integer linear constraints based on (5) for the *star* and *mesh* options. We set $n_{coor} = n_0$ as the *star* central node at the chest, and $N_{hops} = 2$ as the maximum number of re-broadcasting hops in a *mesh* topology. For a two-hop configuration, N_{reTx} is equal to $N^2 - 4N + 5$.

For the application layer, we assume that each node generates 100-byte packets every $1/\phi = 100$ ms by burning $P_{bl} = 100\,\mu$W from a CR2032 coin cell battery. However, the coordinator node (for a *star* topology) relies on larger energy storage to perform its function. Overall, our design space contains 12,288 potential configurations (10 node positions, 3 radio Tx power levels, 2 MAC layer options, and 2 routing schemes), which is too large to be analyzed manually.

4.2 Optimization Results

Figure 3 represents the PDR as a function of the network lifetime (NLT) for the set of feasible configurations suggested by the MILP solver. The feasible configurations span the entire range of PDRs (0 to 100%), and an NLT value from 2 days to more than a month.

The optimal solutions provided by our algorithm for dif-

[1]https://github.com/a-moin/hi-opt.git

[2]Not present in datasheet and based on extrapolation.

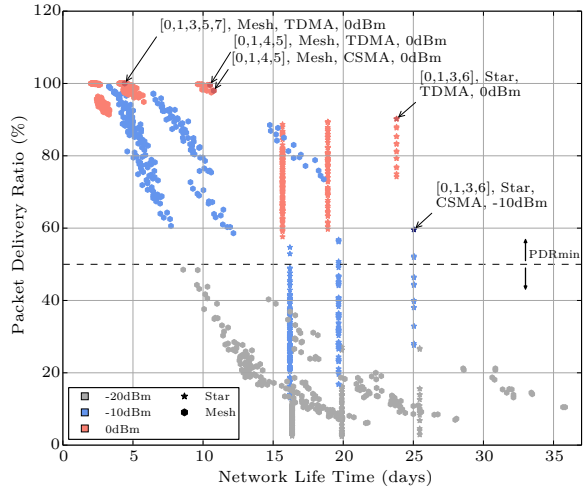

Figure 3: Reliability and lifetime of the feasible network configurations for the optimization problem in Sec. 4.1. The arrows highlight optimal configurations for different values of PDR_{min}.

ferent values of PDR_{min}, all above 50% (as marked by the dashed horizontal line in Figure 3), are shown on the same figure. Each optimization run took on average 52 minutes, resulting into an 87% reduction in the number of required simulations with respect to exhaustive search. In all experiments, Algorithm 1 terminated soon after the first feasible configuration satisfying the PDR constraint was found.

Our algorithm selects a star topology with a minimum number of nodes and a −10-dBm Tx power when the reliability is not a major concern (PDR_{min} below 60%). An increase in Tx power to 0 dBm is suggested for higher reliability requirements, which boosts the PDR up to about 90%. For reliability levels higher than 90%, the routing configuration switches from *star* to *mesh*. This is in agreement with experimental studies [22] showing that a multi-hop (mesh) architecture can indeed produce very high end-to-end PDR at the expense of increased energy consumption. Finally, in a safety-critical application for which 100% reliability is required, a fifth node (n_7 on the shoulder) is added to the mesh to achieve higher redundancy. This, however, comes at the expense of a shorter network lifetime, a couple of days in this case.

The numerical results show that our algorithm enables efficient quantitative evaluation of the design trade-offs across the entire problem space of a Human Intranet and for different topologies. We also compared our approach with simulated annealing [23], a general-purpose method for optimization. On the problem instances reported here, our algorithm runs, on average, 3× faster across the whole range of PDR_{min} values of interest (from 50 to 100%).

5. CONCLUSION

We proposed an optimization-based design exploration approach for a Human Intranet network across the entire communication stack, capable of exploring network lifetime versus reliability trade-offs. Our algorithm combines the solution of a mixed integer linear program to select feasible candidate configurations with accurate discrete-event simulation used to evaluate the performance metrics. Our approach converges faster than general-purpose optimization frameworks on our problem instances, and is able to pinpoint the most promising configurations and network topologies while supporting a broad range of application scenarios.

Acknowledgments The authors wish to acknowledge the faculty and sponsors of the Berkeley Wireless Research Center and the support of TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

6. REFERENCES

[1] J. W. Ng, B. P. Lo, O. Wells, M. Sloman, N. Peters, A. Darzi, C. Toumazou, and G.-Z. Yang, "Ubiquitous monitoring environment for wearable and implantable sensors (UbiMon)," in *Int. Conf. Ubiquitous Computing*, 2004.

[2] B. Gyselinckx, R. J. Vullers, C. Van Hoof, J. Ryckaert, R. F. Yazicioglu, P. Fiorini, and V. Leonov, "Human++: Emerging technology for body area networks." in *VLSI-SoC*, 2006.

[3] J. M. Rabaey, "The human intranet: where swarms and humans meet," in *Proc. Design, Automation and Test in Europe Conf.*, 2015, pp. 637–640.

[4] L. S. Bai, R. P. Dick, P. H. Chou, and P. A. Dinda, "Automated construction of fast and accurate system-level models for wireless sensor networks," in *Proc. Design, Automation and Test in Europe Conf.*, 2011.

[5] I. Beretta, F. Rincon, N. Khaled, P. R. Grassi, V. Rana, and D. Atienza, "Design exploration of energy-performance trade-offs for wireless sensor networks," in *Proc. Design Automation Conf.*, 2012.

[6] *IEEE Std 802.15.4-2006*. IEEE Computer Society, 2006.

[7] P. R. Grassi, I. Beretta, V. Rana, D. Atienza, and D. Sciuto, "Knowledge-based design space exploration of wireless sensor networks," in *Proc. Int. Conf. Hardware/Software Codesign and System Synthesis*, 2012.

[8] P. Nuzzo, H. Xu, N. Ozay, J. B. Finn, A. L. Sangiovanni-Vincentelli, R. M. Murray, A. Donzé, and S. A. Seshia, "A contract-based methodology for aircraft electric power system design," *IEEE Access*, vol. 2, pp. 1–25, 2014.

[9] N. Bajaj, P. Nuzzo, M. Masin, and A. Sangiovanni-Vincentelli, "Optimized selection of reliable and cost-effective cyber-physical system architectures," in *Proc. Design, Automation and Test in Europe Conf.*, 2015.

[10] J. Finn, P. Nuzzo, and A. Sangiovanni-Vincentelli, "A mixed discrete-continuous optimization scheme for cyber-physical system architecture exploration," in *Proc. Int. Conf. Computer-Aided Design*, 2015.

[11] A. Sangiovanni-Vincentelli, "Quo vadis, SLD? Reasoning about the trends and challenges of system level design," *Proceedings of the IEEE*, vol. 95, no. 3, 2007.

[12] D. B. Smith, A. Boulis, and Y. Tselishchev, "Efficient conditional-probability link modeling capturing temporal variations in body area networks," in *Proc. Int. Conf. Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2012.

[13] D. Pediaditakis, Y. Tselishchev, and A. Boulis, "Performance and scalability evaluation of the castalia wireless sensor network simulator," in *Proc. Int. Conf. Simulation Tools and Techniques*, 2010.

[14] *IEEE Standard for Local and Metropolitan Area Networks - Part 15.6: Wireless Body Area Networks*. IEEE Computer Society, 2012.

[15] A. Rahman, W. Olesinski, and P. Gburzynski, "Controlled flooding in wireless ad-hoc networks," in *Int. Workshop on Wireless Ad-Hoc Networks*, 2004.

[16] Y. Khan, A. E. Ostfeld, C. M. Lochner, A. Pierre, and A. C. Arias, "Monitoring of vital signs with flexible and wearable medical devices," *Advanced Materials*, 2015.

[17] www.ibm.com/software/integration/optimization/cplex-optimizer.

[18] https://pypi.python.org/pypi/PuLP.

[19] https://castalia.forge.nicta.com.au/.

[20] http://www.opennicta.com.au/datasets.

[21] http://www.ti.com/lit/gpn/cc2650.

[22] A. Natarajan, B. De Silva, K.-K. Yap, and M. Motani, "To hop or not to hop: Network architecture for body sensor networks," in *Proc. Conf. Sensor, Mesh and Ad Hoc Communications and Networks*, 2009.

[23] https://github.com/perrygeo/simanneal.