

On the Feasibility of Using Wireless Ethernet for Indoor Localization

Andrew M. Ladd, Kostas E. Bekris, Algis P. Rudys, Dan S. Wallach,
and Lydia E. Kavraki

Abstract—IEEE 802.11b wireless Ethernet is becoming the standard for indoor wireless communication. This paper proposes the use of measured signal strength of Ethernet packets as a sensor for a localization system. We demonstrate that off-the-shelf hardware can accurately be used for location sensing and real-time tracking by applying a Bayesian localization framework.

Index Terms—Bayesian inference, sensor fusion, robot localization, wireless Ethernet.

I. INTRODUCTION

The IEEE 802.11b wireless Ethernet standard is becoming increasingly popular and has been deployed in many indoor environments [19]. Many mobile robots already make use of wireless networking for communication. Wireless Ethernet devices measure signal strength as part of their normal operation. We propose the use of off-the-shelf wireless Ethernet adapters on a mobile robot as a tool for global pose estimation. This paper is a feasibility study on the advantages and the difficulties of using this sensor for robot localization. We believe that there is a great potential for wireless Ethernet to be applied as an additional input to a sensor-fusion technique for robot localization.

Determining the pose of the robot from physical sensors is a key problem in robotics, since it plays a pivotal role in various successful mobile robot systems [4]. Outdoor localization can be achieved using global positioning systems (GPS) [18]. Using GPS for indoor localization, however, presents some significant challenges [20]. For the problem of indoor localization, a variety of other sensors have been used, such as vision, infrared (IR) and laser range finders. Mobile robots already employing wireless Ethernet for communication purposes could be retrofitted in software to make use of their adapter as a location sensor. Such a sensor might be very useful for a low-cost robot or team of robots wishing to execute global localization, navigation, and exploration tasks. This is of particular interest for some multirobot configurations; while communicating, the robots could measure signal strengths to each other and engage in collaborative localization.

The chief difficulty in localization with wireless Ethernet is predicting signal strength. Radio frequency (RF) signal strength measured indoors is nonlinear with distance. In addition, it has non-Gaussian noise, resulting from multipath effects and environmental effects, such

Manuscript received November 11, 2002; revised June 16, 2003. This paper was recommended for publication by Associate Editor J. Leonard and Editor S. Hutchinson upon evaluation of the reviewers' comments. The work of A. Ladd was supported in part by FCAR and the National Science Foundation under Grant NSF-IRI-970228. The work of K. Bekris was supported in part by the National Science Foundation under Grant NSF-IRI-970228. The work of A. Rudys and D. Wallach was supported in part by the National Science Foundation under Grant NSF-CCR-9985332, in part by a Texas ATP grant, and in part by a gift from Schlumberger. The work of L. Kavraki was supported in part by the National Science Foundation under Grant NSF-IRI-970228 and Grant NSF-ITR-0205671, and in part by a Sloan Fellowship. This paper was presented in part at MOBICOM, Atlanta, GA, 2002, and in part at IROS, Lausanne, Switzerland, 2002.

The authors are with the Department of Computer Science, Rice University, Houston, TX 77005 USA (e-mail: aladd@cs.rice.edu; bekris@cs.rice.edu; arudys@cs.rice.edu; dwallach@cs.rice.edu; kavraki@cs.rice.edu).

Digital Object Identifier 10.1109/TRA.2004.824948

as building geometry, network traffic, presence of people, and atmospheric conditions.

This paper describes a set of experiments with a laptop carried by a human operator, whereby localization with IEEE 802.11b wireless Ethernet is shown to be feasible. The focus of this work was to determine the usability of wireless Ethernet as a sensor. This was motivated by an application in wireless computing, that of locating an intruder using a laptop [16], [15]. However, it is also a minimalist approach which isolates the sensor we are testing, and generates results applicable to mobile robotics. To compute position, we apply a scheme in the spirit of other Bayesian techniques that have been successfully employed in the context of robotics [23]. We show experiments demonstrating that off-the-shelf wireless hardware can accurately be used for location sensing and tracking with about one-meter precision in a wireless-enabled office building.

1) *Related Work*: The simplest technique used for mobile robot localization has been dead reckoning. With dead reckoning, errors are added to the absolute pose estimate and accumulated. Triangulation techniques were also used. In this case, landmarks are extracted from the sensor input and then they are used to triangulate the robot's position. This works when the sensors are reliable, but leaves several problems unaddressed [6]. Kalman filters were also applied for localization [17], [22], where various sensor data are fused to obtain a new position estimate. This method is provably optimal when noise distributions are Gaussian, but typically fails when this assumption breaks down. The most powerful algorithms to date are based on Bayesian inference, in particular, Markov models [8], [13], and Monte Carlo localization [7], [24]. Most often, the workspace is represented by an occupancy grid.

Alternately, the environment can be modeled with a topological map, e.g., as a generalized Voronoi graph [2]; localization in this paradigm is based on identifying nodes in the graph from geometric environmental information [3], [14]. Our research uses the Bayesian approach. We sample the space, we measure the signal strength at regularly spaced locations, and calculate the probability distribution.

Localization is a problem that has been also explored in the wireless community for a wireless device like a laptop computer, usually carried by a human operator. Many systems have been implemented that use specialized hardware [25], [21]. The RADAR system [1], however, uses only the 802.11b wireless networking for localization and applies nearest-neighbor heuristics and triangulation techniques. The authors report accuracy of 3 m with about 50% probability. While our work has similar design goals to RADAR, we have taken a very different approach.

2) *RF Signal Propagation in Wireless Ethernet*: The IEEE 802.11b standard uses RFs in the 2.4-GHz band, which is license free around the world. Accurate prediction of signal strength from location is a complex and difficult task, since the signal propagates by unpredictable means [5], [19]. In the 2.4-GHz frequency band, microwave ovens, Bluetooth devices, 2.4-GHz cordless phones, and welding equipment can be sources of interference. Signals of this frequency are absorbed by water, and consequently, people will also absorb signal since human bodies are almost 70% water. Due to reflection, refraction, scattering, dependence on atmospheric parameters, and absorption of radio waves by structures inside a building, the transmitted signal most often reaches the receiver by more than one path, resulting in a phenomenon known as multipath fading [11]. Signal multipath effects cause the observed signal strength to vary in unpredictable ways as the receiver position varies, but signal profiles tend to remain approximately the same over short distances [11].

Many efforts have been made to model radio-signal distribution in an indoor environment [9], [19]. Although it has been suggested that the




Fig. 1. Two examples of signal-strength distributions, measured over time at a constant location.

signal propagates according to a log-normal function [9], [10], different experiments have arrived at different distributions, and a general model remains unavailable. Our experiments verified this; we concluded that log-normal fits were only feasible when line-of-sight between transmitter and receiver existed. In our experiments, the noise distributions of signal strength measured at a fixed location varied greatly. In Fig. 1, we show two typical examples of the signal. The two distributions correspond to measurements taken over time from the same position. Although there is a dominant mode in both of them, we observe that the distributions are asymmetric and multimodal (i.e., non-Gaussian).

II. METHODOLOGY

1) *Hardware*: Our experiments were conducted by a human operator carrying a HP OmniBook 6000 laptop with a PCMCIA LinkSys wireless Ethernet card. This particular card uses the Intersil Prism2 chip set. We modified the standard Linux kernel driver to use the base station probe facility of 802.11b [12] to request packets from the base station to obtain the necessary signal-strength measurements. The normal usage of this facility is for determining which base station has the strongest signal and should be chosen as the home station for the card. The card firmware logs and reports eight-bit signal strengths and hardware addresses (MAC addresses) for each response packet received. Each probe returns between zero and four responses from each base station within range of the card. Probes can be made safely at frame rate of between three and six times per second, which will vary with network traffic and location. We remark that this signal is quite thin when compared with other sensors, such as sonar or a laser range finder. We estimate that in our setup, there are roughly five meaningful bits of information from each packet, and these bits are noisy.

2) *Our Model*: The localizer that we implemented operates in the general framework of Bayesian inference localization [8], [13], [23]. We chose a state space and observation space. Position is represented as a probability distribution over the states. The inference calculation consists of conditioning on the observations and then selecting a representative point from the resulting distribution.

We chose points in four hallways of the building where our experiments took place, spaced roughly 1.5 m apart from each other, at two orientations. A point for our experiments was represented as a tuple (x, y, θ) . To summarize, our state space consisted of a set of n points $S = \{s_1 = (x_1, y_1, \theta_1), \dots, s_n = (x_n, y_n, \theta_n)\}$.

Our observation space consisted of the observations that occurred in a single measurement from our base station scanner. A packet consists of k pairs of base station MAC address and signal strength. The number of replies k is different from the number of base stations, which in our experiments is 14. This is due to the fact that some base stations may

not reply at all, or some of them may reply up to four times. We are assuming that the number of times a base station is replying to a probe is independent from the signal strength of the reply. As a result, a single measurement consists of a count k of the number of base station replies, a summary of the frequency counts (the number of times each of the N base stations was seen in this measurement), and then the k pairs. We denote this as a vector $o = \langle k, f_1, \dots, f_N, (b_1, \lambda_1), \dots, (b_k, \lambda_k) \rangle$, where k is the count, N is the total number of base stations, f_i is the frequency count for the i th base station, b_j is the base station index of the j th measurement, and λ_j is the signal strength at that point.

At each point s_i , we take a sample of the observable. For each base station, we build two histograms at that point. The first is a distribution of the frequency counts over the sampled observations. The second is a distribution of signal strengths. Based on this sample, we can calculate two kinds of conditional probability. $\Pr(f_j = a | s_i)$, the probability that the frequency counts for the j th base station is equal to a when we are at state s_i , and $\Pr(\lambda_j | b_j, s_i)$, the probability that the base station b_j has signal strength λ_j at state s_i . For $o = \langle k, f_1, \dots, f_N, (b_1, \lambda_1), \dots, (b_k, \lambda_k) \rangle$, for each $1 \leq i \leq n$ and under the assumption that the signal strength is independent from the number of replies, we compute

$$\Pr(o | s_i) = \left(\prod_{j=1}^N \Pr(f_j | s_i) \right) \cdot \left(\prod_{j=1}^k \Pr(\lambda_j | b_j, s_i) \right).$$

By explicitly integrating a probability distribution of positions based on a received measurement and selecting a representative point, we obtain a position estimate. After trying several possible schemes, we decided to solve a global localization problem for each measurement rather than keep a running estimate, because each measurement usually contains enough information to get a good guess of our position. Initially, we assign equal probability to all possible states in the building of being the actual position. Iterating over all states, we compute which position is most likely to have resulted in the measured signal by selecting the point of maximum probability, assuming that the point exceeds a certain threshold probability. The resulting stream can be further processed to improve precision (see ‘‘Sensor fusion’’ below).

3) *Training Process*: Our system was trained by taking samples every 1.5 m in the world by three different operators. Each operator had to hold a laptop, and to stand still for several seconds at each sample point to collect data. We assume that sampled data was operator independent, that is to say, that we believed that measured distributions would be relatively unaffected by who took the data. The amount of data taken at each point is varied adaptively according to a simple heuristic, which measures the rate of convergence to a stable distribution. Once the sampled distribution at each visible base station had




Fig. 2. Map of Duncan Hall. There are five base stations marked on this map as enumerated small circles (larger circles in the picture are the pillars). Two of them are located at hallway 1 and hallway 2. Overall 14 different base stations were used. They were located on two different floors.

converged beyond a threshold, we halt the process. This allowed us to train the system faster, spending more effort only in positions where it is necessary to achieve an accurate measurement of the signal-strength distribution. In our case, usual sampling times ranged from 10 s to 1 min per position.

4) *Sensor Fusion With a Hidden Markov Model*: We implemented a filter that works on top of the Bayesian inference procedure. It takes the output of the inference engine as a stream of timed observations and tries to stabilize the distribution by noting that a person carrying a laptop typically does not move very quickly. This sort of calculation could be achieved with a much higher degree of precision using odometry from a mobile robot.

We model a moving operator trying to track her position as a hidden Markov model (HMM). We use a more finely discretized state space than the Bayesian inference engine and try to interpolate our position out of the stream of measurements coming from this filter. We observed that, by averaging our training measurements taken at a 1.5 m spacing, and using the average measurement for the points in between, we were able to localize the computer for points where we had not taken any training samples.

For our purposes, an HMM is a set of states $S = \{s_1, \dots, s_n\}$, a set of observations $O = \{o_1, \dots, o_m\}$, a conditional probability $\lambda : S \times O \rightarrow [0, 1]$, and a transition probability matrix A . As in the Bayesian inference engine, each state is a point (x, y, θ) .

The transition probability matrix semantics describe how the system being modeled evolves with time. In this case, it describes how a person travels through the state space. If π is a probability distribution over S , then $\pi' = A\pi$ is the probability distribution after some discrete time step. The idea is that the random state change occurs, “hidden” from the observer.

The observation function λ has semantics identical to observation in the Bayesian inference of position. $\lambda(s, o) = \text{Pr}(o|s)$, the probability of observing o while at s . As each observation arrives, λ is used to update the probability of being in a given state in S , and then A is used to transition states. λ is chosen to be a model of the behavior of the inference engine, and A is chosen to heuristically model human motion.

III. RESULTS

In our experiments, we can measure and track position robustly. In over 70% of our localization attempts, we return either the closest trained position or an adjacent trained position, i.e., we achieved 1.5-m accuracy. Although this may seem as a large error for indoor localization, it must be noted that errors are not accumulated as in the case of odometry, and that this sensor comes almost for free, since we achieved this level of accuracy in a complex indoor environment by employing commodity 802.11b Ethernet equipment. The map of the building we operated in can be found in Fig. 2. It had fairly complicated geometry and the base stations were laid out more than a year before we began our work. We did not make any changes to the original base station locations in our building.

Figs. 3 and 4 show tracking experiments that took place in hallways 1 and 3 of our building, respectively. The operator walked down the hallway, indicating the exact time that certain milestones were passed, giving us an accurate measure of the operator’s true position to compare against the localization results. In the figures, we report both the Bayesian inference static localization results and the HMM fused results. In Fig. 3, a significant improvement is obtained and, overall, the results are excellent. In Fig. 4, the signal was much noisier, due to a relatively poor base station placement. Note that in both cases, errors of 1.5 m are still roughly within one standard deviation.

IV. SUMMARY

Bayesian techniques applied to signal-strength measurements from 802.11b wireless Ethernet allow for real-time and accurate localization. Our work provides a strong indication that both human and robot agents can use their existing wireless network interfaces, which currently serve only as communication devices, as a low-cost localization sensor. The infrastructure for such networks already exists in many real-world environments, and consequently, our scheme can be implemented as a software-only solution.

The experiments were conducted by human operators which introduced error due to signal absorption and lack of odometry. Nevertheless, the results are valid for localization carried with different hardware or with a mobile robot. In fact, similar experiments carried out with a




Fig. 3. Measured error for walking up and down hallway 1, shown on the right graph, is within 1 m with probability 0.64, an improvement of 45% over static localization. This improvement is illustrated in the actual tracking performance, shown in the left graph.

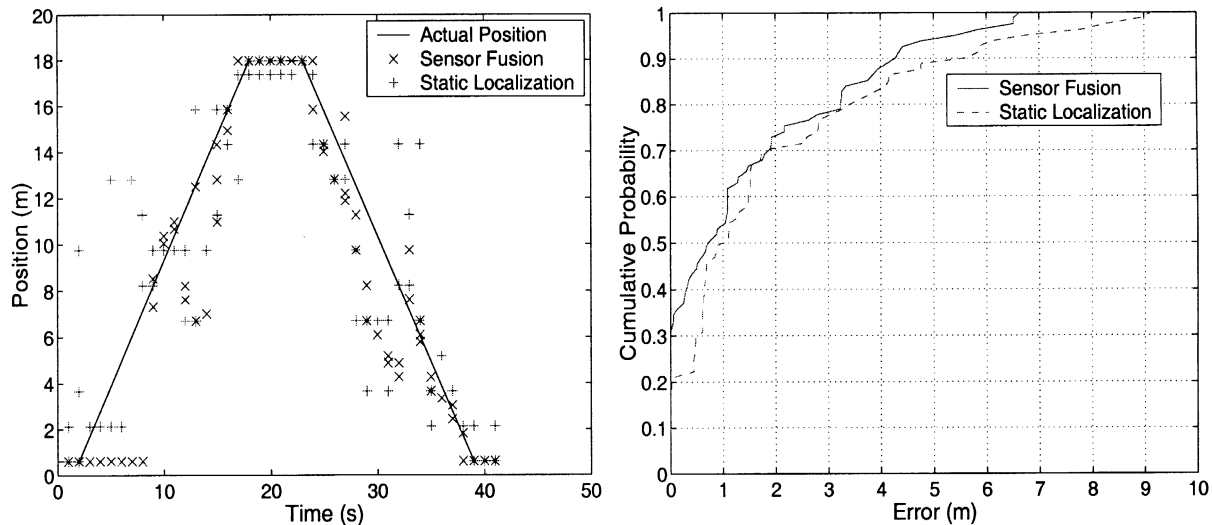


Fig. 4. While sensor fusion provided some improvement, it was not significant, due to pathological behavior of static localization in certain regions, like hallway 3.

robot would likely be significantly more precise, as we avoid absorption errors induced by the operator and have odometry as an additional sensor. Furthermore, there was no effort to improve the quality of the localization procedure by placing the base stations so as to better cover the area where the experiments took place. In an indoor environment, the multipath-fading phenomenon, interference, absorption, reflection, and refraction can not be avoided. However, our results suggest that signal-strength measurement from the 802.11b wireless Ethernet devices are useful cues for localization.

There are many interesting open problems related to the use of wireless Ethernet as a localization sensor, such as the behavior of the signal in dynamic environments. Most of our experiments were taken during the night. During daytime, new problems in localization arise both from the absorption of signal by human bodies and from the heavier network traffic. Furthermore, the framework we have proposed is not restricted to corridor localization where the laptop operator or the robot is forced to walk almost on a straight line. Experiments in large open rooms and on multiple floors can provide further insights on the functionality of wireless Ethernet as a sensor for localization. Last, but not least, is the issue of base station location. Poor initial placement of base stations can

severely decrease the efficiency of localization attempts. Techniques that can provide suggestions for the number and the placement of base stations in order to efficiently cover an environment can significantly improve the performance of localization methods.

ACKNOWLEDGMENT

The authors would like to thank G. Marceau for his valuable suggestions, M. Abdel-Gawad and S. Schell for helping to take measurements, S. Crosby for his comments, and D. Johnson for his advice and comments. Thanks also to the anonymous MOBICOM'02 and IROS'02 reviewers.

REFERENCES

- [1] P. Bahl and V. Padmanabhan, "User Location and Tracking in an In-Building Radio Network," Microsoft Res., Tech. Rep. MSR-TR-99-12, 1999.
- [2] H. Choset and K. Nagatani, "Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization," *IEEE Trans. Robot. Automat.*, vol. 17, pp. 125–137, Apr. 2001.

- [3] H. Choset, S. Walker, K. Eiamsa-Ard, and J. Burdick, "Sensor-based exploration: Incremental construction of the hierarchical generalized Voronoi graph," *Int. J. Robot. Res.*, vol. 19, pp. 126–148, Feb. 2002.
- [4] I. Cox, "Blanche—An experiment in guidance and navigation of an autonomous robot vehicle," *IEEE Trans. Robot. Automat.*, vol. 7, pp. 193–204, Apr. 1991.
- [5] T. Cutler, "Wireless Ethernet and how to use it," in *The Online Industrial Ethernet Book*. Hampshire, U.K.: GGH Marketing Commun., 1999, pt. 5.
- [6] G. Dudek and M. Jenkins, *Computational Principles of Mobile Robotics*. Cambridge, U.K.: Cambridge Univ. Press, 2000.
- [7] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte Carlo localization: Efficient position estimation for mobile robots," in *Proc. 16th Nat. Conf. Artificial Intelligence*, Orlando, FL, 1999, pp. 343–349.
- [8] D. Fox, W. Burgard, and S. Thrun, "Markov localization for mobile robots in dynamic environments," *J. Artif. Intell. Res.*, vol. 11, pp. 391–427, Nov. 1999.
- [9] P. Harley, "Short-distance attenuation measurements at 900 MHz and 1.8 GHz using low antenna heights for microcells," *IEEE J. Select. Areas Commun.*, vol. 7, pp. 5–11, Jan. 1989.
- [10] H. Hashemi, "Impulse response modeling of indoor radio propagation channels," *IEEE J. Select. Areas Commun.*, vol. 11, pp. 967–978, Sept. 1993.
- [11] —, "The indoor radio propagation channel," *Proc. IEEE*, vol. 81, pp. 943–968, July 1993.
- [12] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, ANSI/IEEE Standard 802.11, 1999.
- [13] K. Konolige and K. Chou, "Markov localization using correlation," in *Proc. 17th Int. Joint Conf. Artificial Intelligence*, Seattle, WA, Aug. 1999, pp. 1154–1159.
- [14] B. Kuipers and Y. T. Byan, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations," *J. Robot. Automat. Syst.*, vol. 8, pp. 47–63, 1991.
- [15] A. M. Ladd, K. E. Bekris, G. Marceau, A. Rudys, D. S. Wallach, and L. E. Kavraki, "Using wireless Ethernet for localization," in *Proc. IEEE/RJS Int. Conf. Intelligent Robots and Systems*, vol. 1, Lausanne, Switzerland, Sept.–Oct. 2002, pp. 402–408.
- [16] A. M. Ladd, K. E. Bekris, A. Rudys, G. Marceau, L. E. Kavraki, and D. S. Wallach, "Robotics-based location sensing using wireless Ethernet," in *Proc. 8th ACM Int. Conf. Mobile Computing and Networking*, Atlanta, GA, Sept. 2002, pp. 227–238.
- [17] J. J. Leonard and H. Durrant-Whyte, "Mobile robot localization by tracking geometric beacons," *IEEE Trans. Robot. Automat.*, vol. 7, pp. 376–382, June 1991.
- [18] T. Logsdon, *Understanding the Navstar: GPS, CIS and IVHS*, 2nd ed. New York: Van Nostrand/Reinhold, 1995.
- [19] A. Neskovic, N. Neskovic, and G. Paunovic. (2000, 3rd Quarter) Modern approaches in modeling of mobile radio systems propagation environment. *IEEE Commun. Surveys* [Online]. Available: <http://www.comsoc.org/livepubs/surveys/index.html>
- [20] E. Olsen, C.-W. Park, and J. How, "3D formation flight using differential carrier-phase GPS sensors," in *Proc. ION-GPS Conf.*, Sept. 1998, pp. 1937–1946.
- [21] N. Priyantha, A. Chakraborty, and H. Balakrishnan, "The Cricket location support system," in *Proc. 6th Ann. ACM/IEEE Int. Conf. Mobile Computing and Networking*, Boston, MA, Aug. 2000, pp. 32–43.
- [22] R. Smith and P. Cheeseman, "On the representation of spatial uncertainty," *J. Robot. Res.*, vol. 5, no. 4, pp. 56–68, 1987.
- [23] S. Thrun, "Probabilistic algorithms in robotics," *AI Mag.*, vol. 21, no. 4, pp. 93–109, 2000.
- [24] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, "Robust Monte Carlo localization for mobile robots," *Artif. Intell.*, vol. 101, pp. 99–141, 2000.
- [25] A. Ward, A. Jones, and A. Hopper, "A new location technique for the active office," *IEEE Pers. Commun. Mag.*, vol. 4, pp. 42–47, Oct. 1997.

Integrated Sensing and Filter Design for a Single-Link Flexible Manipulator

G. L. Wang and Y. F. Li

Abstract—This paper addresses the problem of using different sensors in filter design that can simultaneously satisfy multiple specifications. A novel approach is taken in the design paradigm that integrates the sensing strategy with the filter design, which improves the filtering performance. An application to the estimation of the endpoint vibration rate of a single-link flexible manipulator is presented with experimental verifications.

Index Terms—Estimation, filter design, flexible manipulators, performance specifications, sensing.

I. INTRODUCTION

In filter designs, fundamental limitations exist due to the dynamic effects of system structures. Quantifying these limitations is a research topic that has attracted considerable attention [1]. It has been recognized that the sensing strategy and sensing system employed can affect the achievable filtering performance significantly. Classically, the sensing system is predetermined before filter design is conducted. In the past, research in filter design has mainly been concerned with the optimization schemes to include the error variance [2], [3] and \mathcal{H}_∞ -norm minimization [4]. With the optimization-based filter-design method, it is difficult to address a wider range of multiple specifications, owing to the performance constraints.

In this paper, we explore a simple yet efficient way of making use of different sensor resources in filter design. In particular, we propose a design methodology to achieve multiple simultaneous specification (MSS) filtering via an integrated sensing and filtering-design approach. This is achieved by a two-step procedure. The first is to find a finite number of individual specification (IS) filters. This step divides the MSS requirements into several IS design objectives. That is, each IS filter focuses on one or a few specification objectives. The second is the synthesis design stage, where an MSS filter can be found by simply combining the available IS filters in a convex combination. The relevant idea of using MSS for control purposes was studied in [5] and [6]. A new aspect considered here lies in the exploitation of the benefits of potential sensor resources for achieving the MSS filtering. More precisely, the sensing-system selection is incorporated into the complete filter design cycle (see Fig. 1), so as to facilitate the MSS filter design. This is especially useful when rich sensor resources are available.

This paper is organized into two main parts. First, we describe our methodology in the context of bounded-error estimator (BEE) [7]. Second, we present the application of our methodology to a single-link flexible manipulator. The measurement of the flexible modal state can not be directly achieved via hardware sensors. Estimating it using

Manuscript received November 18, 2002. This paper was recommended for publication by Associate Editor C. Melchiorri and Editor A. De Luca upon evaluation of the reviewers' comments. This work was supported by grants from the Research Grants Council of Hong Kong under Project CityU 1049/00E. This paper was presented in part at the American Control Conference, Denver, CO, June 2003.

G. L. Wang was with the Department of Manufacturing Engineering and Engineering Management, City University of Hong Kong, Kowloon, Hong Kong. He is now with Sun Yat-Sen University, Guangzhou, China.

Y. F. Li is with the Department of Manufacturing Engineering and Engineering Management, City University of Hong Kong, Kowloon, Hong Kong (e-mail: meylf@cityu.edu.hk).

Digital Object Identifier 10.1109/TRA.2004.824646