

Distributed Diagnosis of Coupled Mobile Robots

Matthew Daigle, Xenofon Koutsoukos, and Gautam Biswas

Institute for Software Integrated Systems (ISIS)

Department of Electrical Engineering and Computer Science

Vanderbilt University

Nashville, TN 37235

Email: {matthew.j.daigle, xenofon.koutsoukos, gautam.biswas}@vanderbilt.edu

Abstract—Fault diagnosis of coupled mobile robots requires a large number of measurements to be communicated either between the robots or from the robots to a central diagnoser. As computational complexity increases with the number of measurements, centralized algorithms become inefficient. This paper presents a distributed approach for qualitative fault diagnosis of coupled mobile robots. The approach is based on a bond graph modeling framework which incorporates local and distributed control algorithms, multiple sensor types, and both actuator and sensor faults. Relative measurement orderings are introduced to discriminate faults by exploiting the temporal order of the measurement deviations. This increases the discriminatory power of a set of measurements and results in a more efficient qualitative diagnosis algorithm. Distributed diagnosers are designed and applied to coupled mobile robots. Experimental results for a system consisting of two robots pushing a box demonstrate the improvement in both discriminatory power of the measurements and efficiency of the distributed diagnosis approach.

I. INTRODUCTION

Multi-robot teams can be used to autonomously perform a wide range of collaborative tasks in manufacturing, surveillance, and space exploration. In such tasks, detection and diagnosis of faulty behavior is crucial for the system to maintain safe operation. Further, early diagnosis can enable recovery actions. Centralized diagnosis approaches for multi-robot teams are not efficient because they result in a large communication overhead between the robots and do not exploit the computational resources available on each robot.

This paper presents a distributed approach for qualitative fault diagnosis of coupled mobile robots. Our approach is based on the TRANSCEND framework [1], [2] that employs a qualitative approach for analysis of fault transient behavior. This analysis produces fault signatures, which are predicted time-derivative effects of faults on measurements. Distributed diagnosis algorithms based on TRANSCEND are presented in [3].

Our distributed diagnosis approach is based on a bond graph model, which provides a common framework for modeling the physical processes, sensors, and actuators, as well as the communication among the robots. The coupling between the robots produces fault signatures that by themselves do not have the discriminatory power to differentiate between all faults. A new concept, *relative measurement orderings*, is formulated based on the intuition that faults cause deviations in some measurements before others. Relative measurement orderings discriminate faults based on the temporal order of

measurement deviations. A formal diagnosability analysis for single persistent faults shows that using both fault signatures and relative measurement orderings increases the discriminatory power of the measurements and facilitates more efficient diagnoses. Based on this new discriminatory information, distributed diagnosers are designed and applied to coupled mobile robots. Experimental results for a system consisting of two robots pushing a box demonstrate the approach. The results illustrate the advantages of the method, namely increasing the discriminatory power of the measurements, and improving the efficiency of the distributed diagnosis approach.

The paper is organized as follows. Section II describes the problem and presents the model used for diagnosis. Section III presents the diagnosis architecture, and Section IV addresses distributed fault detection. Section V discusses the fault isolation approach. Section VI demonstrates the approach using experimental results for a system of two mobile robots pushing a box. Section VII discusses related work and Section VIII concludes the paper.

II. PROBLEM FORMULATION

Our diagnosis framework applies to mobile robots that collaborate to perform a task. In this paper, we focus on a system consisting of two robots simultaneously pushing a box. The control objective of the system is to push the box along a straight line perpendicular to a wall at a predefined velocity, v_{ref} , keeping the box edge parallel to the wall, as depicted in Figure 1.

Each robot includes a local controller that regulates the velocities of its wheels. The sensor suite includes motor encoders to measure wheel velocity, a gyroscope to measure heading, and a laser range finder to measure distance to the

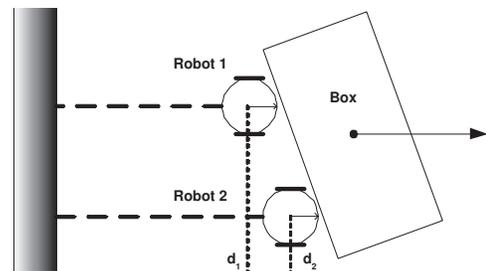


Fig. 1. Multi-robot box-pushing setup

wall. A distributed controller coordinates the collaborative task by determining the desired velocity for each robot based on local and remote sensor measurements, communicated via a wireless network.

A. Modeling of a Single Mobile Robot

Each robot is modeled using a bond graph model similar to the model presented in [4]. Bond graphs are topological models of dynamical systems that are particularly suitable for diagnosis because they incorporate both causal and temporal information [5]. We extend the model presented in [4] by modeling the wheels of each robot with local controllers, the sensors, and sensor faults.

The bond graph for a single robot is given in Figure 2. Bonds (energy transfer pathways) are represented as half arrows and signals (information transfer pathways) as arrows. The robot plant consists of the left wheel, right wheel, and body subsystems. Inertia components model wheel mass and inertia m_w , robot mass M_c , and rotational inertia J_c . The capacitor components (with parameters C_L and C_R) model the mechanical stiffness of the robot, and the resistor elements (with parameters R_L and R_R) model energy dissipation in the system. Transformers model the transformations between linear and rotational velocities. The 1-junctions represent the common velocity points: ω_L , the rotational velocity of the left wheel, ω_R , the rotational velocity of the right wheel, v , the forward velocity of the robot, and ω , the rotational velocity of the robot. The 0-junctions represent common force points on the left and right sides of the robot, F_L and F_R .

Sensor models in the bond graph are derived from the kinematic relationships between the robot velocities and the measurements. The laser range finder and gyroscope sensors use the kinematic equations based on the linear and rotational velocities of the robot body. Robot i 's position from the wall, $d_i(t)$, and heading, $\theta_i(t)$, are described by:

$$\begin{aligned} \dot{d}_i(t) &= v_i(t) \cos \theta_i(t) \\ \dot{\theta}_i(t) &= \omega_i(t). \end{aligned}$$

The equations for the optical encoder measurements involve a gain transforming the wheels' rotational velocities to their linear velocities, i.e., $v_{L,i}(t) = G_{EL} \omega_{L,i}(t)$ and $v_{R,i}(t) = G_{ER} \omega_{R,i}(t)$, where for Robot i , $\omega_{L,i}$ and $\omega_{R,i}$ are the rotational velocities of the left and right wheels, and G_{EL} and G_{ER} are the encoder gains for the left and right wheels, respectively.

The sensors are modeled in the bond graph as modulated sources of flow that encapsulate these equations. For the gyroscope, the flow source is the rotational velocity of the robot, ω . In the bond graph, this is represented by the flow variable f_{20} , associated with bond 20. The measured variable, the heading, is e_{29} (the effort variable associated with bond 29), which is the integral of ω plus the sensor bias (if any). The laser range finder model includes a gyroscope model that provides the true heading to the flow source for the laser component. The flow equation is $f_9 \cos e_{38}$ (or $v \cos \theta_0$, where θ_0 is the true heading). This value gets integrated to produce

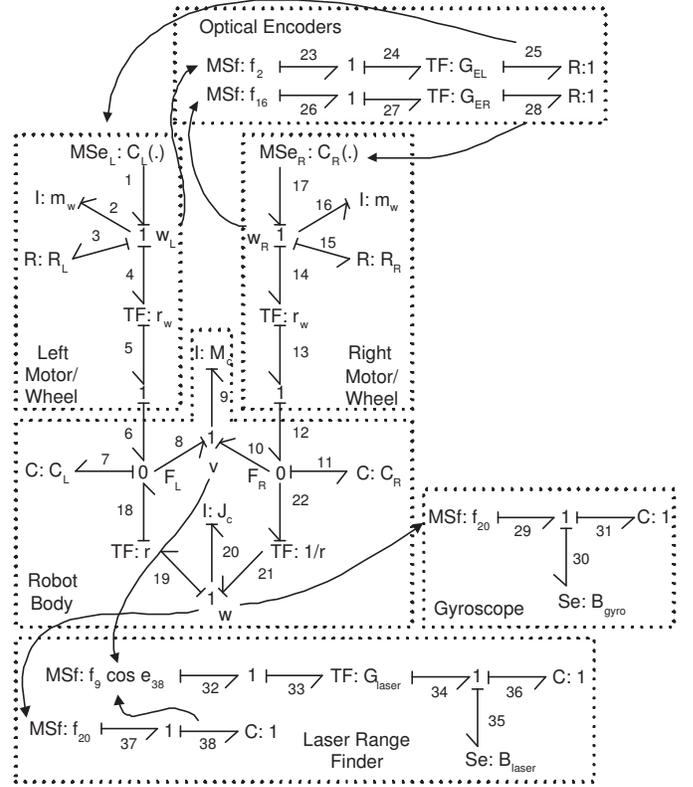


Fig. 2. Bond graph model of a single robot

the distance measurement. The measured value is then e_{34} , which is the true distance reading plus the sensor bias (if any). For the case of the optical encoders, the flow is the rotational velocity of a wheel (ω_L and ω_R) passed through a gain, so the measured variables are f_{25} and f_{28} .

Local PID controllers are also modeled in the bond graph. The inputs to the robots, the motor torques, are modeled as modulated sources of effort which encapsulate the wheel control equations. The torque for the left (right) wheel is represented in the bond graph by the modulated source of effort MSe_L (MSe_R). The PID controller is represented by the function $C_L(\cdot)$ ($C_R(\cdot)$) that modulates the torque. The edges from the observed velocities to the wheel sources represent the control loops of the PID controllers.

B. Modeling of Coupled Mobile Robots

We have implemented a simple distributed control scheme to solve the box-pushing problem based on the protocols presented in [6]. Since the control scheme is made explicit in the bond graph model, our approach can be used with any other control scheme. We have implemented a distributed controller whose objective is to compute the reference velocities of the wheels, $v_{dL,i}$ and $v_{dR,i}$, based on the measurements. The distributed controllers communicate with each other at a fixed rate.

Each robot tries to move forward at the desired speed, v_{ref} , while keeping even with the other robot. To achieve this, each robot communicates its measured position. The forward

velocity for Robot i is computed by

$$v_{d,i}(k) = v_{ref} + g_d(d_j(k) - d_i(k)), \quad (1)$$

where g_d is a gain selected such that the system is stable, and d_j is the distance measurement of Robot j . The desired heading is $\theta = 0$, so the desired rotational velocity is given by another proportional control law:

$$\omega_{d,i}(k) = -g_\omega \theta_i(k), \quad (2)$$

where g_ω is an appropriate gain. Equations (1) and (2) can be decoupled into individual left and right wheel velocities:

$$\begin{aligned} v_{dL,i}(k) &= v_{d,i}(k) - r \omega_{d,i}(k) \\ v_{dR,i}(k) &= v_{d,i}(k) + r \omega_{d,i}(k), \end{aligned}$$

where r is the radius of the robot. Therefore, the distributed controller computes the reference velocities

$$\begin{aligned} v_{dL,i}(k) &= v_{ref} + g_d(d_j(k) - d_i(k)) + r g_\omega \theta_i(k) \\ v_{dR,i}(k) &= v_{ref} + g_d(d_j(k) - d_i(k)) - r g_\omega \theta_i(k) \end{aligned}$$

for the wheels to satisfy the control objective.

The distributed control is modeled in the bond graph in the same manner as the local PID control. Signals are introduced from relevant measurements to the wheel sources, including from the remote distance measurement representing the communication. The multi-robot bond graph is the composition of two single robot bond graphs with these signal edges and the distributed control functions included.

C. Modeling for Diagnosis

Faults are represented as abrupt parameter value changes in the bond graph model. Actuator (motor) faults are modeled as changes in the effort sources. A saturation fault in an actuator limits the maximum wheel velocity. Actuator failure can be viewed as a saturation fault with a limit of zero. Sensor bias is modeled as an additive fault, and is represented by a change in the effort source at the measured value (nominally, the effort is 0). For example, a bias in the laser range finder manifests as an abrupt, constant value added to the true measurement value. Sensor failures are modeled as multiplicative faults, and are parameterized by a change in the sensors' transformer gains. For the optical encoders, the nominal value of G_{EL} (or G_{ER}) is r_w , the wheel radius, and in failure the gain is in the interval $[0, r_w)$, i.e., some percentage of the encoder counts are missed (at least 20%). For the laser, the nominal transformer gain is 1, and in failure the gain becomes 0, thus the measured value does not change after the point of failure. Table I shows the mapping of faults to parameter changes in the bond graph model (a superscript of + or - indicates the direction of change of the parameter value).

III. DIAGNOSIS APPROACH

The diagnosis architecture for the multi-robot system consists of four core components as illustrated in Figure 3. The local observers, implemented as Kalman filters, are based on a state space model of each robot derived from the bond graph.

TABLE I
MAPPING OF FAULTS TO PARAMETER CHANGES

Fault	Parameter
Left Actuator Saturation/Failure	MSe_L^-
Right Actuator Saturation/Failure	MSe_R^-
Left Encoder Saturation/Failure	G_{EL}^-
Right Encoder Saturation/Failure	G_{ER}^-
Gyroscope Bias	B_{gyro}^+, B_{gyro}^-
Laser Range Finder Bias	B_{laser}^+, B_{laser}^-
Laser Range Finder Failure	G_{laser}^-

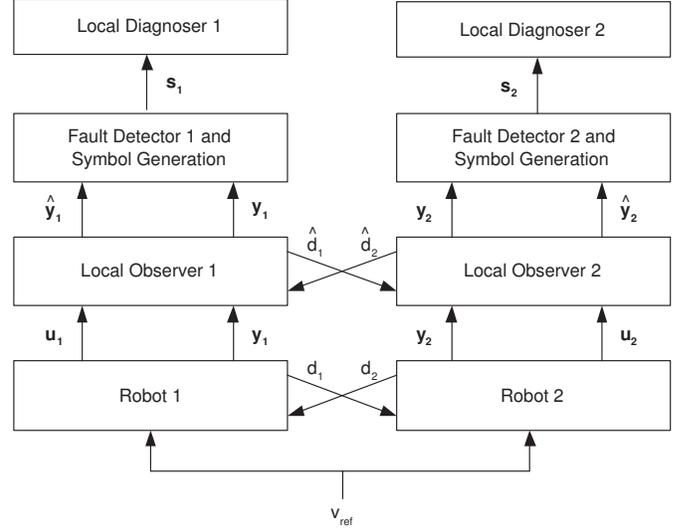


Fig. 3. Diagnosis architecture

They compute the output estimates, \hat{y}_i , given the input u_i , the local measurements y_i , and the position measurement and its estimate for the remote robot (Robot j), i.e., d_j and \hat{d}_j . The fault detectors compute the residuals of the measurements as the difference between actual and predicted values. If a fault is detected, the symbol generator computes qualitative values, s_i (fault signatures), for the changes in measurement values. Each local diagnoser uses these signatures to isolate the fault.

IV. DISTRIBUTED FAULT DETECTION

Residual-based fault detection methods utilize a model of the system to compute the residual. For an ideal system with a perfect model, any nonzero residual indicates a fault. Because of noise and model imperfections, the measured and predicted values may differ under nominal conditions. Therefore, a Kalman filter is used to track the system trajectory, and only statistically significant deviations from zero indicate a fault. This reduces the false alarm rate significantly.

In order to employ this fault detection strategy, we use a distributed, decentralized, extended Kalman filter (DDEKF) [7]. This method creates local filters for each robot which communicate relevant observations and estimates to the other robot. The local state vectors for each robot must contain all state variables needed to produce estimates of the ob-

served variables. Each robot observes its own wheel velocities, heading, and wall distance, i.e., the local measurements are $y_i = [d_i \ v_{Li} \ v_{Ri} \ \theta_i]^T$ for Robot i . State space equations can be directly derived from the bond graph model [5]. Unknown parameters can be identified using system identification techniques. A discretized, reduced order form of the identified state space model is used, assuming the dynamics of the wheels are decoupled. For the reduced model, the local state vector for Robot i is $x_i = [d_i \ r_1^1 \ r_1^2 \ r_1^3 \ r_1^4 \ l_1^1 \ l_1^2 \ l_1^3 \ l_1^4 \ \theta_i \ d_j]^T$, where j is the remote robot, and r_i^h and l_i^h correspond to dynamic states of the left and right wheels, respectively.

The difference between the observed values and the estimated values define the residual. A statistical test can then be used to detect a fault. In this paper, we use the Z-test [8] on the residuals to determine, given the estimated variance of the residuals, a confidence interval, and the modeling error, whether a fault has occurred. A small sliding window of samples is used to estimate the current mean of the residuals, and this is preceded by a much larger sliding window to estimate the variance. When the current mean of one of the residual signals shows a statistically significant deviation from zero (accounting for modeling error), a fault is detected.

V. FAULT ISOLATION

A. Background

The TRANSCEND architecture [1], [2] is employed for diagnosis of the multi-robot system. Fault isolation in TRANSCEND is based on a qualitative analysis of the transient dynamics caused by abrupt faults. Deviations in measurement values after a fault occurrence constitute a fault signature, where predicted deviations in magnitude and higher order derivative values are mapped to symbols of the set $\{+, 0, -\}$, which correspond to deviations above normal, no deviations, and deviations below normal, respectively.

Fault isolation in TRANSCEND utilizes a Temporal Causal Graph (TCG) representation, which can be derived directly from the bond graph model of the system. The TCG captures the causal and temporal relations between system variables. It specifies the signal flow graph of the system in a form where edges are labeled with single component parameter values or direct or inverse proportionality relations. Figure 4 depicts the TCG model for a single robot, with state variables circled and measured variables boxed. The TCG of the entire system consists of the two TCGs for the robots, with additional edges that convey the distance measurement of one robot to the effort sources of the other robot. This captures the qualitative effects of the remote distance measurements in the transient dynamics of the robot's motion.

Fault signatures are generated by running a forward-propagation algorithm on the TCG to predict qualitative effects of faults on measurements [1]. The qualitative effect of a fault, $+$ or $-$, is propagated to all measurement vertices in the TCG to determine fault signatures for each measurement. It can be shown that these provide a temporal progression of the predicted qualitative changes in the measured signal. By

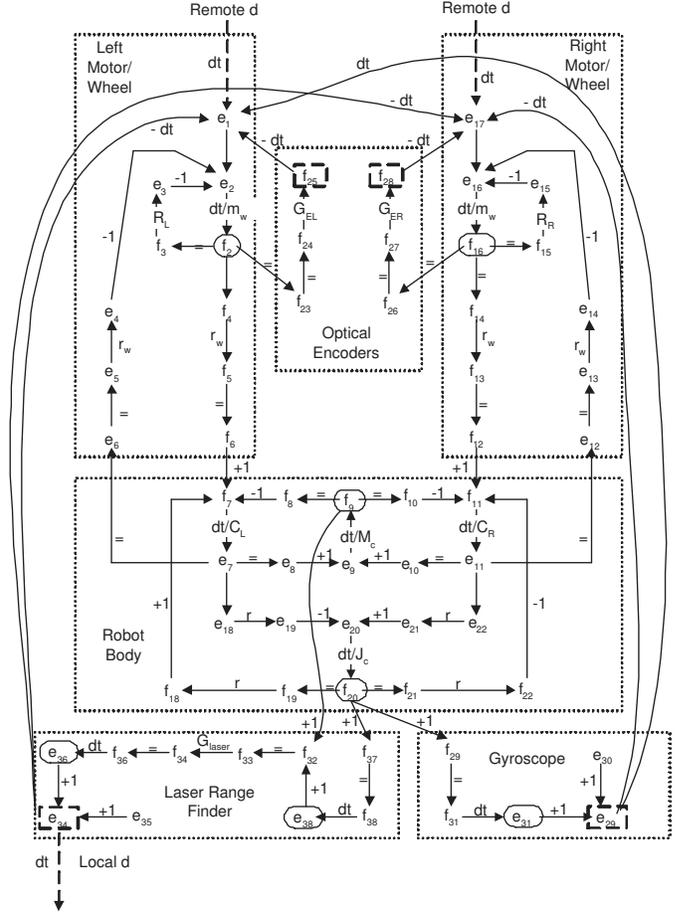


Fig. 4. TCG for a single robot of the multi-robot system

expressing the fault signature as derivative effects, measurement analysis can be formulated as a progressive monitoring scheme, where lower order changes manifest before higher order changes.

In the robot TCG, for example, the fault MSe_L^- starts at the vertex e_1 . The $-$ change propagates to the distance measurement (vertex e_{34}) by passing through four temporal edges ($e_2 \rightarrow f_2$, $f_7 \rightarrow e_7$, $e_9 \rightarrow f_9$, and $f_{36} \rightarrow e_{36}$) without the sign getting reversed, thus the first change is a 4th order change of $-$. This change will eventually manifest as a change in magnitude and slope, which can be reliably measured.

Fault isolation in TRANSCEND compares measurement residual magnitude and slopes to predicted fault signatures. Fault hypotheses whose signatures are consistent with the measured residual symbols are retained, and others are dropped. Diagnosis in TRANSCEND is based only on fault signatures as the discriminating information.

B. Diagnosability Analysis

Proper design of a diagnoser requires the system to be diagnosable, i.e. all faults of interest can be uniquely isolated with the given measurement set. Table II shows the fault signatures for the multi-robot system generated from the system TCG, with only the magnitude change symbol and the first non-zero

TABLE II

FAULT SIGNATURES FOR MULTI-ROBOT BOX-PUSHING SYSTEM, WITH
LEADING ZEROS REMOVED

Fault	$v_{L,1}$	$v_{R,1}$	θ_1	d_1	$v_{L,2}$	$v_{R,2}$	θ_2	d_2
$MS_{L,1}^-$	0-	0*	0+	0-	0-	0-	0*	0-
$MS_{R,1}^-$	0*	0-	0-	0-	0-	0-	0*	0-
$G_{EL,1}^-$	-+	0*	0-	0+	0+	0+	0*	0+
$G_{ER,1}^-$	0*	-+	0+	0+	0+	0+	0*	0+
$B_{gyro,1}^+$	0+	0-	+ -	0*	0*	0*	0*	0*
$B_{gyro,1}^-$	0-	0+	-+	0*	0*	0*	0*	0*
$B_{laser,1}^+$	0-	0-	0*	+ -	0+	0+	0*	0+
$B_{laser,1}^-$	0+	0+	0*	-+	0-	0-	0*	0-
$G_{laser,1}^+$	0+	0+	0*	0-	0-	0-	0*	0-
$MS_{L,2}^-$	0-	0-	0*	0-	0-	0*	0+	0-
$MS_{R,2}^-$	0-	0-	0*	0-	0*	0-	0-	0-
$G_{EL,2}^-$	0+	0+	0*	0+	-+	0*	0-	0+
$G_{ER,2}^-$	0+	0+	0*	0+	0*	-+	0+	0+
$B_{gyro,2}^+$	0*	0*	0*	0*	0+	0-	+ -	0*
$B_{gyro,2}^-$	0*	0*	0*	0*	0-	0+	-+	0*
$B_{laser,2}^+$	0+	0+	0*	0+	0-	0-	0*	+ -
$B_{laser,2}^-$	0-	0-	0*	0-	0+	0+	0*	-+
$G_{laser,2}^+$	0-	0-	0*	0-	0+	0+	0*	0-

direction of change symbol shown. A * symbol indicates an indeterminate effect, i.e., there are at least two paths of the same order that propagate + and - effects, and the dominant effect is unknown. From the signatures, it is clear that an actuator fault on one robot cannot be distinguished from an actuator fault on another robot, because the signatures for both faults are consistent for all measurements. Therefore, using the given measurement set and the fault signatures approach, the system is not globally diagnosable. This motivates the need for employing additional discriminatory information to achieve global diagnosability.

C. Relative Measurement Orderings

Relative measurement orderings refer to the intuition that fault effects will manifest in some parts of the system before others. For example, a fault occurring in one robot will likely manifest first in that robot and then in the remote robot, if there are energy storage elements in the path between the local and remote sensors in the bond graph. If there are no energy storage elements, the relation is algebraic and no delay will be observed.

Definition 1: Consider a fault f and measurements m_1 and m_2 ; if the fault manifests in m_1 before m_2 then we can define a relative measurement ordering between m_1 and m_2 for fault f , denoted as $m_1 \prec_f m_2$.

Relative measurement orderings can be derived from the TCG based on the notion of a *fault path*.

Definition 2: A fault path for a fault f and measurement m is a path in the TCG which begins at the fault f and ends at the measurement m .

The set of all fault paths from f to m is denoted by $FP_{f,m}$. The order of a fault path is defined as the number of temporal

edges in the path. A minimum order fault path is a path in $FP_{f,m}$ that contains the minimum number of temporal edges needed to reach m from f . More than one fault path of a specific order may exist for f and m , since there are often multiple paths from one vertex to another in the TCG.

Definition 3: The minimum order fault path set for f and m is the set of all minimum order fault paths, and is denoted as $FP_{f,m}^*$.

A fault path represents the temporal propagation of a fault to a specific measurement variable in the system. For a certain fault, there are multiple fault paths leading to a measurement. Since lower order effects of faults manifest themselves first [2], only the minimum order fault path sets are useful in determining relative measurement orderings. For this purpose, we define a method of comparing fault paths.

Definition 4: For $p \in FP_{f,m_1}$ and $p' \in FP_{f,m_2}$, p is a temporal subpath of p' ($p \sqsubset p'$) if all temporal edges in p exist in p' in the same ordering, and the order of p is less than the order of p' .

Theorem 1: If for every $p' \in FP_{f,m_2}^*$ there exists $p \in FP_{f,m_1}^*$ such that $p \sqsubset p'$, then we have $m_1 \prec_f m_2$.

Proof: In the signal flow graph for the TCG, let r_1 be the measurement vertex corresponding to m_1 , r_2 the vertex for m_2 , and r_f the successor vertex of the edge with fault parameter f . The transfer functions from r_f to r_1 , $R_1(s)$ and from r_f to r_2 , $R_2(s)$, can be derived. Assume for every $p' \in FP_{f,m_2}^*$ there exists $p \in FP_{f,m_1}^*$ such that $p \sqsubset p'$. Then each minimum order path from r_f to r_2 must go through r_1 or a vertex which can be expressed as $r_1 \cdot G$, where G is some constant gain. $R_2(s)$ is a sum of terms which each correspond to different forward paths from r_f to r_2 . Because lower order effects manifest first, terms that correspond to forward paths of non-minimum order can be removed to produce $R_2'(s)$. Similarly, $R_1'(s)$ can be produced. Because every minimum order path from r_f to r_2 goes through a vertex $r_1 \cdot G$, $R_1'(s)$ must appear as a factor in each term of $R_2'(s)$, therefore $R_2'(s) = H(s)R_1'(s)$, where $H(s)$ is a proper transfer function. The order of m_1 is less than the order of m_2 by the definition of the \sqsubset relationship, so the number of poles for $R_1'(s)$ must be less than the number for $R_2'(s)$. $H(s)$ must introduce more poles than zeros to $R_2'(s)$, and, therefore, $H(s)$ is strictly proper. From $H(s)$, we can discretize using the given sampling rate of the system to get $H(z)$. Since $H(s)$ is strictly proper, $H(z)$ is, therefore $r_2'(k) = f(r_1'(k-1))$. Since $r_2'(k)$ depends only on past values of $r_1'(k)$, with appropriately selected detection thresholds¹, a deviation resulting from fault f will appear first in m_1 and then in m_2 , thus $m_1 \prec_f m_2$. ■

Therefore, for a given fault, we can say that it manifests in measurement m_1 before measurement m_2 if for all minimum order fault paths of m_2 , there is a minimum order fault path for m_1 the fault will traverse before completely traversing the given fault path of m_2 . The transient due to the fault is slower

¹This guarantees that for some time $|r_1(k)|$ will be greater than $|r_2(k)|$, after that time $|r_2(k)|$ may overtake $|r_1(k)|$ depending on the gain of $H(z)$. Therefore thresholds must be small enough such that deviations will cross them before that time.

for m_2 than for m_1 , thus, the fault will manifest first in m_1 and then in m_2 . If this ordering is violated, we can eliminate the fault hypothesis.

For example, consider an actuator fault of the left wheel of Robot 1, $MSe_{L,1}^-$. The minimum order fault path set for the velocity measurement of Robot 1, $v_{L,1}$, consists of the path $\{e_1 \rightarrow e_2 \rightarrow f_2 \rightarrow f_{23} \rightarrow f_{24} \rightarrow f_{25}\}$, which contains only one temporal edge with label dt/m_w , implying an integration effect. Minimum order fault path sets for all other measurements must pass through that same edge, thus the temporal subpath relation holds. Therefore, we can define the ordering $v_{L,1} \prec_{MSe_{L,1}^-} m$ for all other measurements m .

Definition 5: An ordering set for a fault f , R_f , is the set of all relative measurement orderings for fault f .

Definition 6: A conflict between ordering sets R_{f_1} and R_{f_2} for measurement set M exists if there are two measurements $m_i, m_j \in M$ such that $\{m_i \prec_{f_1} m_j\} \in R_{f_1}$ and $\{m_j \prec_{f_2} m_i\} \in R_{f_2}$.

For a given measurement set and for each fault, we can derive a set of fault signatures and also a set of measurement orderings from the TCG. Signatures alone have been used to distinguish between different faults in [1], [2]. However, the ordering sets can also be used as further distinguishing information for fault isolation. Therefore, the discriminatory power of a set of measurements is enhanced by using both fault signatures and relative measurement orderings. For a given set of measurements, two faults can be discriminated if they have different fault signatures or if they have conflicts in their ordering sets. Further, these two notions are independent and can be combined to distinguish among fault hypotheses.

Using this information, actuator faults can now be globally distinguished. From the global TCG model, it follows that an actuator fault will appear first in the velocity measurement of that wheel, and then in other measurements. Therefore, if an actuator fault occurs in Robot 1, it will detect the fault before Robot 2, and vice versa. Thus actuator faults occurring on different robots can be distinguished using relative measurement orderings.

Analysis of the multi-robot system shows that faults manifest first in their associated measurement before other measurements in the system (e.g., actuator and encoder faults manifest first in velocity measurements of that wheel). More importantly, since the only connection between the robots is in the distance measurement signal, local faults manifest first in a local measurement, and remote faults manifest first in a remote measurement. Therefore, if a local measurement deviates before a remote measurement, the fault must be local.

D. Distributed Diagnosis

If the system is globally diagnosable, then a distributed approach can improve the efficiency of the diagnosis. Each local diagnoser isolates faults in its subsystem using local measurements and some remote measurements, if required. Since accessing remote measurements is expensive, our design goal is to find the minimum number of remote measurements (0 or more) that make each subsystem locally diagnosable.

Algorithm 1 Distributed Diagnoser Design

```

Input: local fault sets  $F_i$ , local measurement sets  $M_i$ , fault signatures, ordering sets,  $k$  subsystems
for subsystem  $i \in 1, \dots, k$  do
  identify set  $F'_i \subseteq F_i$  such that  $f \in F'_i$  cannot be completely distinguished using  $M_i$ 
  for  $f \in F'_i$  do
    identify minimum set of communicated measurements to globally diagnose  $f$ 
    add this set to the local measurement set
  end for
end for

```

The design approach ensures that a local diagnosis will be globally correct. Assuming single persistent faults, since the local diagnosers achieve a global diagnosis, this avoids the need for a centralized coordinator. A distributed algorithm has been developed using only fault signatures in [3]. In this paper, we extend the algorithm to incorporate relative measurement orderings as additional distinguishing information.

The algorithm generates the distributed diagnoser by minimizing the number of shared measurements between subsystems. For each subsystem, if a fault is not globally diagnosable using local measurements, it searches neighboring subsystems for a minimal set of additional measurements to make the fault globally diagnosable. The pseudocode is given as Algorithm 1. In the worst case all combinations of measurements are considered, so the algorithm is exponential. Practically, since the diagnosers are built offline, their design time complexity is not of much concern.

For the box-pushing system, the subsystems are the individual robots. The diagnoser for Robot i is responsible for diagnosing the faults $\{MSe_{L,i}^-, MSe_{R,i}^-, G_{EL,i}^-, G_{ER,i}^-, B_{gyro,i}^+, B_{gyro,i}^-, B_{laser,i}^+, B_{laser,i}^-, G_{laser,i}^-\}$ using measurements $\{v_{L,i}, v_{R,i}, \theta_i, d_i\}$, i.e., each robot is responsible for diagnosing local faults using its local measurements.

Running the algorithm shows that communicating only the distance measurements between the robots is enough to achieve the design goal, i.e., the diagnoser of Robot 1 needs $\{v_{L,1}, v_{R,1}, \theta_1, d_1, d_2\}$ as its measurement set, and the diagnoser of Robot 2 needs $\{v_{L,2}, v_{R,2}, \theta_2, d_2, d_1\}$. Since the residual values for the distance measurements are already (indirectly) communicated for fault detection, there is no added communication cost. Also, by using predicted measurement orderings for the distance measurements, the diagnosers implicitly eliminate all remote faults if a local measurement deviates before a remote measurement. Therefore, as soon as one robot isolates a single fault hypothesis, the fault must be local and a global diagnosis is known. This occurs because from the global diagnosis model, we know that locally, all local faults appear first in a local measurement, and all remote faults appear first in the communicated remote measurement.

Each local diagnoser runs an online fault isolation algorithm [1]. The algorithm starts with the set of local fault candidates and their associated fault signatures after an initial deviation has been detected. It matches the candidates' predicted fault

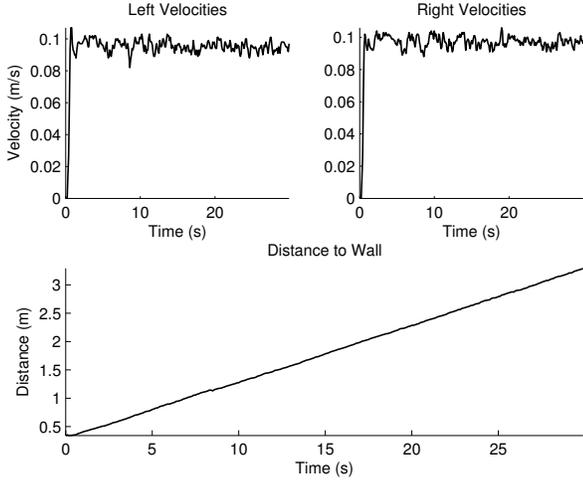


Fig. 5. Nominal trajectory for Robot 1

signatures to observed measurement deviations as they appear, dropping candidates whose signatures are inconsistent with observed transients. This algorithm is augmented such that candidates are also dropped if there is an inconsistency between predicted measurement orderings and observed measurement orderings. Using this new information also makes fault isolation more efficient, because less measurements are required to uniquely isolate a fault, and the knowledge that a certain measurement has not yet deviated is helpful. Using only fault signatures, that measurement must deviate before it can be used to discriminate faults.

VI. EXPERIMENTAL RESULTS

The distributed detection and diagnosis algorithms were demonstrated with two ActivMedia Pioneer 3-DX mobile robots communicating over an 802.11b wireless ad-hoc network. All faults listed in Table I were introduced through software. The sampling period of the distributed controllers and diagnosers was 0.1 seconds. At the selected sampling rate, the packet loss was negligible (measured less than 0.1%).

Figure 5 shows the nominal trajectory for Robot 1 pushing a box at a desired speed of 0.1 m/s. Robot 2 has a similar trajectory. The robots were able to achieve the desired speed and their distance error stayed small (less than 20 mm).

In the following, we illustrate our approach for a complete failure of the right wheel of Robot 1 ($MSe_{R,1}^-$). Figure 6 shows the measurements for the robots, and Table III traces the diagnosis steps. Initially, the diagnosers assume empty fault sets. The fault is injected at $t = 15.0$ seconds. The fault causes the right wheel to slow down, therefore, the left wheel slows to maintain the heading. Robot 2 slows down to keep the box parallel to the wall. A deviation in $v_{R,1}$ at 15.3 seconds triggers Robot 1's fault isolation procedure, and it starts with its entire fault set, F_1 , as its set of possible candidates. As predicted, $v_{R,1}$ is the first deviation, so based on orderings, the fault set is reduced to only faults of the right wheel. The change of $v_{R,1}$ matches the fault signature of 0-, thus isolating the fault

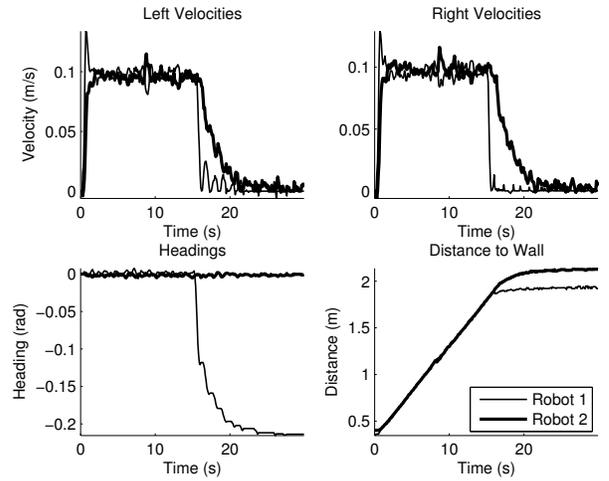


Fig. 6. Trajectories for a right actuator failure on Robot 1

TABLE III
DIAGNOSIS TRACE FOR RIGHT ACTUATOR FAULT OF ROBOT 1

Time	Event	Fault set, Robot 1	Fault set, Robot 2
15.0	Fault injected (1)	\emptyset	\emptyset
15.3	Fault detected (1) $v_{R,1}$ deviates 0- observed Diagnosis (1)	F_1 $\{MSe_{R,1}^-, G_{ER,1}^-\}$ $\{MSe_{R,1}^-\}$	\emptyset
15.4	θ_1 deviates	$\{MSe_{R,1}^-\}$	\emptyset
15.5	$v_{L,1}$ deviates	$\{MSe_{R,1}^-\}$	\emptyset
15.7	Fault detected (2) d_1 deviates Diagnosis (2)	$\{MSe_{R,1}^-\}$	F_2 \emptyset
16.6	$v_{L,2}$ deviates $v_{R,2}$ deviates	$\{MSe_{R,1}^-\}$	\emptyset \emptyset
18.2	d_2 deviates	$\{MSe_{R,1}^-\}$	\emptyset

to be $MSe_{R,1}^-$. By design, this is guaranteed to be the globally unique fault, so Robot 2 can be notified and recovery actions may commence. Only one measurement needed to deviate to obtain a global diagnosis, so this demonstrates the efficiency of using relative measurement orderings in fault isolation.

For further illustration, we allow Robot 2 to go on without notification of the diagnosis result. At 15.7 seconds, d_1 deviates, triggering the diagnoser of Robot 2. Its diagnoser begins with its entire fault set, F_2 . The change of d_1 matches a 0-. Using only signatures, it is not guaranteed that the observed effects were not caused by a remote fault, as shown in Section V. In this case, Robot 2 observes the remote measurement d_1 deviate before any of its local measurements, so it can eliminate all of its faults, i.e., the fault must be remote.

All faults listed in Table I were successfully isolated using the distributed diagnoser. The summary of the diagnosis results is shown in Table IV. The degree of the fault and its time of injection are shown, along with all measurements deviations observed until a global diagnosis is known. All times are in seconds.

TABLE IV
DIAGNOSIS RESULTS

Fault	Injection time	Detection and diagnosis time
$MSe_{L,1}^-$ (Failure)	15.0	15.3, $v_{L,1}$ 0-
$MSe_{L,1}^-$ (Sat. 0.05 m/s)	15.0	15.3, $v_{L,1}$ 0-
$MSe_{L,1}^-$ (Sat. 0.08 m/s)	15.0	15.5, $v_{L,1}$ 0-
$MSe_{R,1}^-$ (Failure)	15.0	15.3, $v_{R,1}$ 0-
$MSe_{R,1}^-$ (Sat. 0.05 m/s)	15.0	15.3, $v_{R,1}$ 0-
$MSe_{R,1}^-$ (Sat. 0.08 m/s)	15.0	15.5, $v_{R,1}$ 0-
$G_{EL,1}^-$ (Failure)	15.0	15.0, $v_{L,1}$ - *
$G_{EL,1}^-$ (50% decrease)	15.0	15.0, $v_{L,1}$ - *
$G_{EL,1}^-$ (20% decrease)	15.0	15.1, $v_{L,1}$ - *
$G_{ER,1}^-$ (Failure)	15.0	15.0, $v_{R,1}$ - *
$G_{ER,1}^-$ (50% decrease)	15.0	15.0, $v_{L,1}$ - *
$G_{ER,1}^-$ (20% decrease)	15.0	15.1, $v_{L,1}$ - *
$B_{gyro,1}^+$ (3 degrees)	15.0	15.1, θ_1 + *
$B_{gyro,1}^+$ (5 degrees)	15.0	15.0, θ_1 - *
$B_{laser,1}^+$ (40 mm)	15.0	15.1, d_1 + *
$B_{laser,1}^-$ (40 mm)	15.0	15.1, d_1 - *
$G_{laser,1}^-$	15.0	15.3, d_1 0-

VII. RELATED WORK

Fault detection and isolation in single mobile robot systems has been addressed previously. In [4], a mobile robot was modeled using a simplified bond graph model and actuator faults were diagnosed using the fault signature approach. In [9], the specific problem of fault detection was addressed by developing a technique which accounted for both kinematic and dynamic behaviors in order to generate better residuals in spite of parametric uncertainty. Work in [10] addressed sensor fault detection and identification in a mobile robot by using a bank of Kalman filters, each modeling the robot under a different fault condition. Identification was accomplished by using probabilistic methods. This work was extended in [11] by using a neural network to detect and identify both sensor and mechanical failures based on the output of the filter bank. Particle filtering techniques have also been employed in diagnosis of robots in [12], [13].

In this paper, a single, distributed Kalman filter is used to produce estimates, and the residuals are analyzed to determine qualitative fault transient behavior for fault isolation. To our knowledge, this is the first time a distributed approach is developed and demonstrated for coupled mobile robots.

VIII. CONCLUSIONS

In this paper we presented an approach for distributed diagnosis of coupled mobile robots. We presented a system model encompassing the plant, sensors, actuators, communication, and control. The DDEKF was applied for the purpose of distributed fault detection. A qualitative fault analysis archi-

tecture, TRANSCEND, was extended to incorporate new diagnostic information, relative measurement orderings, to achieve global system diagnosability. Relative measurement orderings increase the discriminatory power of a set of measurements. Therefore diagnosers can require fewer measurements, and diagnoses are achieved faster. Distributed diagnosers were designed based on a global system model using this new information. The design was such that they achieved independent, global diagnoses. Experimental results demonstrated the validity and usefulness of the approach.

Future work will address further generalizing this approach to larger multi-robot systems, such as formations of robots. In this case, the analysis becomes more difficult because of the more complex robot interactions, but the underlying modeling framework can still be used.

ACKNOWLEDGMENT

This work was supported in part by NSF CNS-0452067 and NSF CNS-0347440.

REFERENCES

- [1] P. Mosterman and G. Biswas, "Diagnosis of continuous valued systems in transient operating regions," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 29, no. 6, pp. 554–565, 1999.
- [2] E.-J. Manders, S. Narasimhan, G. Biswas, and P. Mosterman, "A combined qualitative/quantitative approach for fault isolation in continuous dynamic systems," in *SafeProcess 2000*, vol. 1, Budapest, Hungary, June 2000, pp. 1074–1079.
- [3] I. Roychoudhury, G. Biswas, X. Koutsoukos, and S. Abdelwahed, "Designing distributed diagnosers for complex physical systems," in *Proceedings of the 16th International Workshop on Principles of Diagnosis (DX 05)*, Monterey, California, June 2005, pp. 31–36.
- [4] M. Ji, Z. Zhang, G. Biswas, and N. Sarkar, "Hybrid fault adaptive control of a wheeled mobile robot," *IEEE/ASME Transactions on Mechatronics*, vol. 8, no. 2, pp. 226–233, June 2003.
- [5] D. C. Karnopp, D. L. Margolis, and R. C. Rosenberg, *Systems Dynamics: Modeling and Simulation of Mechatronic Systems*, 3rd ed. New York: John Wiley & Sons, Inc., 2000.
- [6] B. R. Donald, J. Jennings, and D. Rus, "Information invariants for distributed manipulation," *International Journal of Robotics Research*, vol. 16, no. 5, pp. 673–702, 1997.
- [7] A. G. Mutambara, *Decentralized Estimation and Control for Multisensor Systems*. Boca Raton: CRC Press, 1998.
- [8] R. E. Kirk, *Statistics: An Introduction*. Fort Worth: Harcourt Brace, 1999.
- [9] W. Dixon, I. Walker, and D. Dawson, "Fault detection for wheeled mobile robots with parametric uncertainty," in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, vol. 2, July 2001, pp. 1245–1250.
- [10] S. Roumeliotis, G. Sukhatme, and G. Bekey, "Sensor fault detection and identification in a mobile robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, Victoria, BC, Canada, Oct 1998, pp. 1383–1388.
- [11] P. Goel, G. Dedeoglu, S. Roumeliotis, and G. Sukhatme, "Fault detection and identification in a mobile robot using multiple model estimation and neural network," in *IEEE International Conference on Robotics and Automation*, vol. 3, 2000, pp. 2302–2309.
- [12] V. Verma, G. Gordon, R. Simmons, and S. Thrun, "Real-time fault diagnosis," *IEEE Robotics & Automation Magazine*, vol. 11, no. 2, pp. 56–66, June 2004.
- [13] R. Dearden and D. Clancy, "Particle filters for real-time fault detection in planetary rovers," in *Proceedings of the International Workshop on Principles of Diagnosis (DX'02)*, Semmering, Austria, 2002, pp. 1–6.