Best-first Enumeration Based on Bounding Conflicts, and its Application to Large-scale Hybrid Estimation (Extended Abstract)*

Eric Timmons† and **Brian C. Williams**
Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology
{etimmons, williams}@mit.edu

Abstract

State estimation methods based on hybrid discrete and continuous state models have emerged as a method of precisely computing belief states for real world systems, however they have difficulty scaling to systems with more than a handful of components. Classical, consistency based diagnosis methods scale to this level by combining best-first enumeration and conflict-directed search. While best-first methods have been developed for hybrid estimation, conflict-directed methods have thus far been elusive as conflicts summarize constraint violations, but probabilistic hybrid estimation is relatively unconstrained. In this paper we present an approach (A*BC) that unifies best-first enumeration and conflict-directed search in relatively unconstrained problems through the concept of "bounding" conflicts, an extension of conflicts that represent tighter bounds on the cost of regions of the search space. Experiments show that an A*BC powered state estimator produces estimates up to an order of magnitude faster than the current state of the art, particularly on large systems.

1 Introduction

There is a continuously growing demand for complex systems with autonomous decision making capabilities that are robust and safe. These desires can be achieved using systems that have the ability to self-repair by using planners online to generate novel responses to exceptional situations. A key capability needed by such systems is the ability to accurately estimate the system state. While discrete models have long been a mainstay of the model-based reasoning community [De Kleer and Williams, 1987], these models do not have the requisite resolution needed when controlling or detecting incipient failures in dynamic hybrid discrete and continuous systems.

While exact hybrid estimation is theoretically simple — given an appropriate continuous state estimator, generate a continuous state estimate for every possible discrete state trajectory — it quickly becomes infeasible due to the exponential growth in discrete trajectories over time. As such, the current state of the art in hybrid state estimation focuses on approximate estimation. These techniques include Multiple Model (MM) methods such as the Generalized Pseudo-Bayesian Algorithm (GPB) [Ackerson and Fu, 1970], the detection-estimation method [Tugnait, 1982], the residual correlation Kalman filter bank [Hanlon and Maybeck, 2000], the Interacting Multiple Model (IMM) algorithm [Blom and Bar-Shalom, 1988], and adaptive MM methods by Li et al. [1996; 1999; 2000]. More recently, techniques such as the Hybrid Mode Estimator (HME) [Hofbaur and Williams, 2002; Hofbaur and Williams, 2004], the Hybrid Diagnostic Engine (HyDE) [Narasimhan and Brownston, 2007], and combined stochastic and greedy estimation [Blackmore *et al.*, 2008] have also been developed.

While these state of the art techniques have been shown to effectively estimate the hybrid state of small subsystems comprised of a hand-full of components, they have difficulties scaling to larger, real-world systems. Consistency based state estimators from the model-based reasoning community are able to scale in large part through the use of best-first enumeration and conflict-directed search [Williams and Ragno, 2007], as well as stochastic search methods [Feldman *et al.*, 2010]. Conflict-directed search serves the role of efficiently pruning large sets of inconsistent states (i.e., states with zero probability), best-first enumeration focuses the estimator on the states that are most likely, and stochastic methods allow the algorithms to remove the burden of completeness. While all methods are important, conflicts have been shown to be particularly effective due to their pruning ability. While scaling of hybrid estimation methods has been improved through best-first [Hofbaur and Williams, 2002] and sampling-based methods [Blackmore *et al.*, 2008], the creation of effective conflict-directed methods has proven more challenging. The primary difficulty is that a conflict is a consistency based concept that represents sets of states that have zero probability based on a proof of logical inconsistency. However, in the extreme case of a stochastic environment with unbounded uncertainty (such as Gaussian noise models), all behaviors are consistent, albeit unlikely.

In this paper we present an approach to hybrid estimation that augments best-first enumeration with a variant of

*This paper is an extended abstract of an article in the Journal of Artificial Intelligence Research [Timmons and Williams, 2020].

†Contact Author

conflicts that is more suitable to probabilistic, rather than consistency-based, inference through the concept of *bounding conflicts*, as well as a search algorithm (A* with Bounding Conflicts) that uses them. Instead of representing a set of search states that are inconsistent, bounding conflicts compactly encode an area of the search space where the predicted cost — such as one obtained from a heuristic bounding function used during search — is much lower than the true cost, as well as a tighter bounding function to use in those areas of the state space. Bounding conflicts are similar in spirit to *valued nogoods* [Dago and Verfaillie, 1996], but more powerful as bounding conflicts provide a tighter bound function instead of a static bound.

While the focus of this paper is on hybrid estimation, we gain insight into bounding conflicts by viewing hybrid estimation as an instance of an optimization problem. As such, we present both a general best-first enumeration algorithm based on bounding conflicts, and a hybrid estimation method based on this general capability. For hybrid estimation, this allows the search for mode assignments to learn which modes are unlikely given the observations and avoid them to quickly focus in on the best candidates. However, these poor mode assignments aren't discarded, instead they are saved for later, to be expanded and tested if there is enough time.

In the remainder of this paper we describe the hybrid discrete and continuous state estimation problem as well as how the systems are modeled. Then we outline a best-first hybrid state estimation approach and how it can be solved as an instance of tree search. Next we introduce the A* with Bounding Conflicts (A*BC) algorithm and discuss how bounding conflicts can be learned in hybrid state estimation. Last, we provide empirical evidence showing our approach produces state estimates up to an order of magnitude faster than the state of the art.

2 Problem Statement

Hybrid state estimation is an instance of state filtering for systems with both discrete and continuous state. The goal of a hybrid state estimator is to compute a belief over the system state at time t ($\mathbf{x}_t$) given a model of the system, a starting belief, observations of the system ($\mathbf{y}_{1:t}$), and the control inputs ($\mathbf{u}_{1:t}$). The state of the system is fully determined by discrete mode variables ($\mathbf{m}_t$) and continuous state variables ($\mathbf{x}_{c,t}$), resulting in Equation 1 which describes the belief state.

$$p(\mathbf{m}_t, \mathbf{x}_{c,t} | \mathbf{y}_{1:t}, \mathbf{u}_{1:t}) \tag{1}$$

We model the state space and dynamics of the system being estimated using a Concurrent Probabilistic Hybrid Automaton (CPHA). CPHAs consist of a number of Probabilistic Hybrid Automata (PHAs) operating concurrently, interacting via constraints on shared variables. This composition makes CPHAs particularly useful for modeling large scale systems composed of many individual components operating in concert.

In the CPHA modeling formalism, the discrete state of a system or component is called its *mode*. The mode of the system determines the evolution of both the continuous and discrete state of the system between successive time steps. This is modeled by having the mode of a component determine which algebraic and differential equations are active during a given time step. The relationships between the various variables at successive time steps are visualized as a dynamic Bayes net in Figure 1.

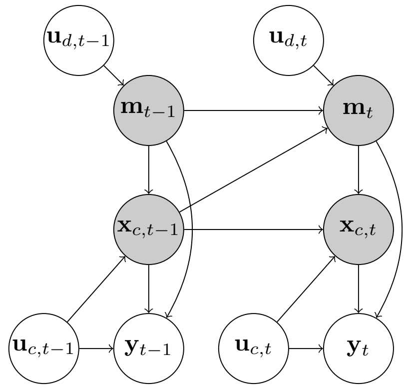

Figure 1: The recursive hybrid estimation problem shown as a dynamic Bayes net. The hidden variables are shaded.

Given an assignment to every mode variable at time t, the system wide continuous dynamics and outputs can be computed, typically with a symbolic solver, as equations of the form below. Equation 2 represents the evolution of the system-wide continuous state and Equation 3 represents the continuous observations of the system.

$$\mathbf{x}_{c,t} = \mathbf{f}_t(\mathbf{x}_{c,t-1}, \mathbf{u}_{c,t}, \mathbf{v}_{s,t}) \tag{2}$$

$$\mathbf{y}_{c,t} = \mathbf{g}_t(\mathbf{x}_{c,t}, \mathbf{u}_{c,t}, \mathbf{v}_{o,t}) \tag{3}$$

3 Approach

Due to the hybrid discrete and continuous nature of the state, we represent the belief state is as a mixture of N independent sub-beliefs[1]. The i'th sub-belief at time t ($\hat{\mathbf{x}}_t^{(i)}$) is parameterized by a weight ($w_t^{(i)}$), mode assignment ($\hat{\mathbf{m}}_t^{(i)}$), and continuous probability distribution ($p_{c,t}^{(i)}$). Intuitively, a sub-belief states that with probability $w_t^{(i)}$, the system is in the specified mode and the belief state over the continuous variables is described by the continuous distribution. The entire belief state can be constructed using a weighted sum over the sub-beliefs.

Given this representation of the belief state, Hofbaur and Williams [2004] have shown that a sub-belief $\hat{\mathbf{x}}_t^{(j)}$ can be recursively computed from a sub-belief at the previous time step $\hat{\mathbf{x}}_{t-1}^{(i)}$ and a new mode assignment $\hat{\mathbf{m}}_t^{(j)}$. The weight computation is shown in Equations 4 and 5 where P_T and P_O are the *hybrid transition and observation likelihoods*. Computing ($p_{c,t}^{(i)}$) is then accomplished using any continuous state estimator; for this work, we use Kalman filter variants such as the Extended Kalman Filter (EKF) [Sorenson, 1985], Unscented Kalman Filter (UKF) [Julier and Uhlmann, 1997], and Truncated Unscented Kalman Filter [Teixeira *et*

[1]In mixture models, these are normally called the mixture's components. We choose to use "sub-beliefs" to eliminate any confusion with the components that make up the system being estimated.

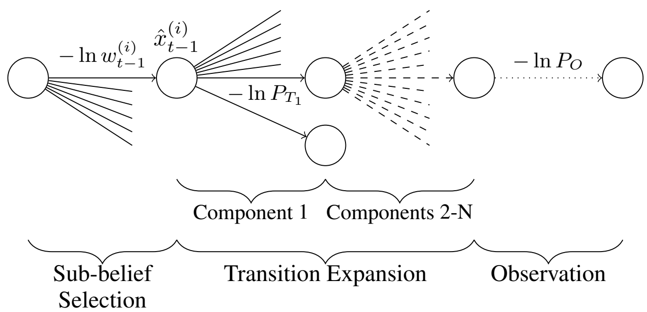

Figure 2: Search tree for hybrid state estimation with a fixed variable order.

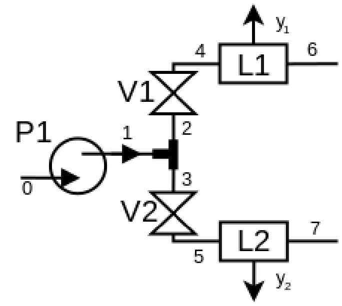

Figure 3: A simplified fluid cooling system. The loads L1 and L2 are producing heat, have their temperatures observed, and can be cooled by fluid being pumped by P1. Valves V1 and V2 can be opened and closed to direct the flow.

al., 2010; Simon, 2010; Garcia-Fernandez *et al.*, 2012].

$$w_{t|t-1}^{(j)} = P_T(\hat{\mathbf{m}}_t^{(j)}, \hat{\mathbf{x}}_{t-1}^{(i)}, \mathbf{u}_{d,t}) w_{t-1}^{(i)} \quad (4)$$

$$w_t^{(j)} = \eta P_O(\mathbf{y}_t, p_{c,t}^{(j)}, \mathbf{u}_{c,t}) w_{t|t-1}^{(j)} \quad (5)$$

The difficulty in computing a complete belief state is that the number of mode assignments is exponential in the number of components in the system. Additionally, each sub-belief is dependent on the mode trajectory and the number of mode trajectories grows exponentially over time. As such, for any non trivial system, the belief state must be approximated. We approximate the belief state using the k sub-beliefs with the highest weights. In order to determine the k best sub-beliefs at a given time, we frame it as a decision problem: given $\hat{\mathbf{x}}_{t-1}$, determine the pairs of sub-beliefs from $t-1$ and successor modes that result in the largest weights.

We propose solving this decision problem using best-first search, namely a variant of A* called A* with Bounding Conflicts (A*BC). A search tree with a fixed variable ordering is visualized in Figure 2. The first layer of the tree represents choosing a sub-belief from time $t-1$ and the cost of the choice is the negative log of that sub-belief's weight. The next N layers of the tree represent choosing the modes of the N components of the system, the combined cost of these decisions is the negative log of the hybrid transition likelihood. The last layer represents the incorporation of observations into the system and its cost is the negative log of the hybrid observation likelihood. The goal is to find the k paths with the lowest cost.

It is straightforward to compute a tight heuristic to guide A* search in the first $N+1$ layers of the tree: for every component not yet assigned a mode, assume that it takes the most likely mode transition. However, the hybrid observation likelihood is a function of the continuous belief state at time t, which makes computing a tight heuristic difficult. This results in a potentially very loose bound of 0 cost being used in practice. A cost of 0 for the last layer is realized only when the actual observations perfectly match the most likely observations. This weak heuristic can result in the search being "tricked" into exploring large amounts of the state space (the paths with the highest *a priori* probabilities) even if the observations are very unlikely in those modes.

Instead, when A*BC discovers a subpath where the heuristic predicts a cost to go that is much lower than the actual cost, it learns a *bounding conflict* (defined below) that summarizes why the heuristic was poor and uses that bounding conflict to compute better heuristics in the future. In the case of hybrid state estimation, bounding conflicts are learned when P_O is evaluated (using the results from the EKF continuous state estimator) and shows that the observations are unlikely ($P_O < 0.5$).

Definition 1. *A bounding conflict is a pair $\langle z, b \rangle$ where:*

- z is a partial assignment to the decision variables of a best-first enumeration problem and
- $b : P_z \to \mathbb{R}$ is a function that maps extensions of z to a bound on the cost of any further extensions.

For an example bounding conflict, consider the simplified fluid system pictured in Figure 3. Assume that the system model says there is a 90% probability of valve V1 being closed. Additionally, assume that the state estimator has observed the temperature of L1 has increased less than would be expected if there were no coolant flow. When the hybrid state estimator computes P_O for this scenario and notes the discrepancy between expected and true observations, it can learn the bounding conflict $\langle V_1 = \text{closed}, b_1 \rangle$ where b_1 is a function that captures the factor by which P_O can be reduced in any mode that contains V_1 is closed.

A*BC then uses these bounding conflicts to both compute a tighter bound for any given node in the search tree and dynamically change the order in which it searches to proactively steer away from sub spaces in the tree where the cost is known to be high. This is accomplished in large part by Algorithm 1. SPLIT-ON-BOUNDING-CONFLICT is one of A*BC's two methods to generate the neighbors for a search node. Given a bounding conflict γ, the method generates a set of neighbors using two techniques. The latter technique (lines 2-8) is the classical conflict directed search technique; every neighbor generated this way is guaranteed to be inconsistent with γ's partial assignment. The remaining neighbor (line 1) generated is simply the partial mode assignment of the search node, extended to include γ's partial assignment. However, this neighbor is also annotated to include the information that it manifests the bounding conflict γ, allowing A*BC to use the tighter bounding function contained within to push the node deeper into the search queue and further expand it only if necessary.

Algorithm 1: SPLIT-ON-BOUNDING-CONFLICT

Input: The node to exapnd, n, and the conflict to split on, γ.
Output: The children of n that resolve the conflict.

children ← MAKE-NODE($\{\gamma[z]\}, n$);
foreach assignment $\in \gamma[z]$ **do**
  $x \leftarrow$ variable of assignment;
  $y \leftarrow$ value of assignment;
  foreach $v \in (dom(x) - y)$ **do**
    children ← MAKE-NODE($\{n[z] \cup \{x = v\}\}, n$);
  end
end
return children;

In order to learn bounding conflicts in the hybrid state estimation domain, we propose an approach that augments the derivation of the system's continuous state dynamics for a given mode (Equations 2 and 3) with causal analysis [Nayak, 1995; Trave-Massuyes and Pons, 1997]. This analysis allows the state estimator to work backward from observations that are far from their expected values, to the equations that were used to produce those expected values, and finally to the mode assignments that activated those equations. This allows the state estimator to nearly trivially generate bounding conflicts using the output from the Kalman filter based continuous estimator.

4 Results

To demonstrate the effectiveness of the A*BC algorithm as a best-first enumerator for hybrid state estimation, it was implemented and compared to HME without bounding conflicts on two example systems. The first is a reproduction of the three PHA system from Hofbaur [2004]. The second is an analog for a shipboard cooling system consisting of sixty non-trivial components (a mixture of pumps, valves, loads, check valves, and flow meters).

Direct comparison to HME without A*BC as the search algorithm on the simple three PHA system shows that using A*BC results in approximately 25% fewer derivations and runs of Kalman filters and a corresponding approximate 25% reduction in runtime for a variety of values of k. Previous work has already shown that HME without A*BC is faster than alternative methods such as IMM on this problem.

The gains of using the A*BC search algorithm become more apparent on larger systems. For the cooling system test, we ran the proposed A*BC approach and the A* based HME approach for a fixed period of time and recorded the number of sub-beliefs the estimator was able to prove had the highest weights as a function of time. Typical results are summarized in Figure 4.

These results show that the proposed approach generates sub-beliefs at the same rate as HME in the beginning, but is quickly able to generate sub-beliefs faster than HME. This allows the proposed A*BC approach to generate more sub-beliefs in a fixed period of time, meaning that its estimates cover more probability mass in its approximation. This results in the proposed approach being less likely to prune the correct mode assignment.

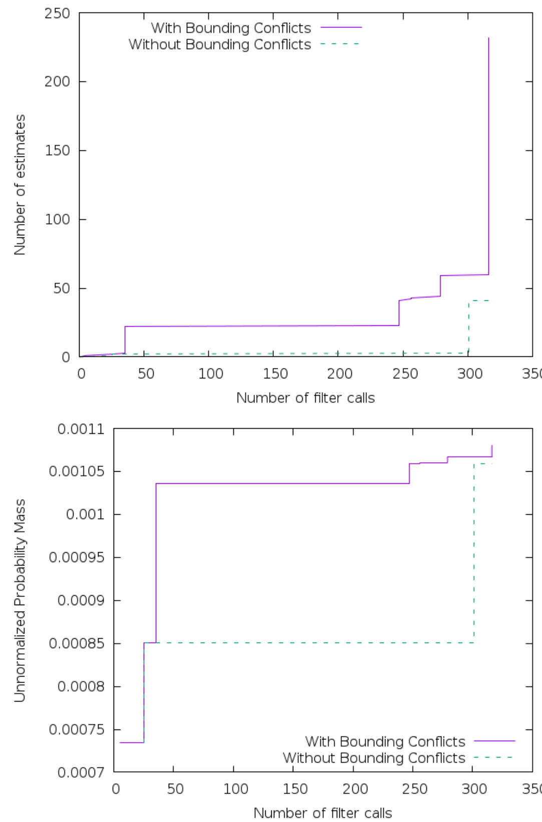

Figure 4: Typical results of A*BC vs. A* on a fluid system analog. The top graph shows the number of sub-beliefs produced as a function of Kalman filter invocations (a proxy for time). The bottom graph shows the (unnormalized) probability mass covered.

5 Conclusion

In this paper, we have introduced *bounding conflicts*, a novel extension of conflicts that describe both where a search algorithm's bounding function is not tight and a tighter bounding function for that region of the state space. Additionally, we have provided a best-first enumeration algorithm based on bounding conflicts (A*BC) and have described a state estimator for hybrid discrete and continuous systems built on top of this enumerator.

This new state estimator using A*BC outperforms the previous state of the art estimator for large-scale hybrid systems. It produces the best sub-beliefs with fewer Kalman filter executions, allowing a better state estimate to be produced in less time.

Acknowledgments

This work was funded by the Office of Naval Research through the Johns Hopkins University Applied Physics Laboratory and by the NASA Planetary Science and Technology through Analog Research (PSTAR) program.

References

[Ackerson and Fu, 1970] G Ackerson and K Fu. On state estimation in switching environments. *IEEE Transactions on Automatic Control*, 15(1):10–17, 1970.

[Blackmore *et al.*, 2008] Lars Blackmore, Stanislav Funiak, and Brian C. Williams. A combined stochastic and greedy hybrid estimation capability for concurrent hybrid models with autonomous mode transitions. *Journal of Robotic and Autonomous Systems*, 56(2):105–129, February 2008.

[Blom and Bar-Shalom, 1988] Henk AP Blom and Yaakov Bar-Shalom. The interacting multiple model algorithm for systems with markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780–783, 1988.

[Dago and Verfaillie, 1996] Pierre Dago and Gérard Verfaillie. Nogood recording for valued constraint satisfaction problems. In *Proceedings Eighth IEEE International Conference on Tools with Artificial Intelligence*, pages 132–139. IEEE, 1996.

[De Kleer and Williams, 1987] Johan De Kleer and Brian C Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32(1):97–130, 1987.

[Feldman *et al.*, 2010] Alexander Feldman, Gregory Provan, and Arjan Van Gemund. Approximate model-based diagnosis using greedy stochastic search. *Journal of Artificial Intelligence Research*, 38:371–413, 2010.

[Garcia-Fernandez *et al.*, 2012] Ángel F Garcia-Fernandez, Mark R Morelande, and Jesús Grajal. Truncated unscented kalman filtering. *IEEE Transactions on Signal Processing*, 60(7):3372–3386, 2012.

[Hanlon and Maybeck, 2000] Peter D Hanlon and Peter S Maybeck. Multiple-model adaptive estimation using a residual correlation kalman filter bank. *IEEE Transactions on Aerospace and Electronic Systems*, 36(2):393–406, 2000.

[Hofbaur and Williams, 2002] Michael W Hofbaur and Brian C Williams. Mode estimation of probabilistic hybrid systems. In *Hybrid Systems: Computation and Control*, pages 253–266. Springer, 2002.

[Hofbaur and Williams, 2004] Michael W Hofbaur and Brian C Williams. Hybrid estimation of complex systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 34(5):2178–2191, 2004.

[Julier and Uhlmann, 1997] Simon J Julier and Jeffrey K Uhlmann. New extension of the kalman filter to nonlinear systems. In *Signal Processing, Sensor Fusion, and Target Recognition VI*, volume 3068, pages 182–193. International Society for Optics and Photonics, 1997.

[Li and Bar-Shalom, 1996] Xiao-Rong Li and Yaakov Bar-Shalom. Multiple-model estimation with variable structure. *IEEE Transactions on Automatic control*, 41(4):478–493, 1996.

[Li *et al.*, 1999] X Rong Li, Xiaorong Zwi, and Youmin Zwang. Multiple-model estimation with variable structure. iii. model-group switching algorithm. *IEEE Transactions on Aerospace and Electronic Systems*, 35(1):225–241, 1999.

[Li, 2000] X Rong Li. Multiple-model estimation with variable structure. ii. model-set adaptation. *IEEE Transactions on Automatic Control*, 45(11):2047–2060, 2000.

[Narasimhan and Brownston, 2007] Sriram Narasimhan and Lee Brownston. Hyde–a general framework for stochastic and hybrid model-based diagnosis. *Proc. DX*, 7:162–169, 2007.

[Nayak, 1995] Pandurang Nayak. *Automated Modeling of Physical Systems*. Lecture Notes in Artificial Intelligence. Springer, 1995.

[Simon, 2010] Dan Simon. Kalman filtering with state constraints: a survey of linear and nonlinear algorithms. *IET Control Theory & Applications*, 4(8):1303–1318, 2010.

[Sorenson, 1985] Harold Wayne Sorenson. *Kalman filtering: theory and application*. IEEE, 1985.

[Teixeira *et al.*, 2010] Bruno OS Teixeira, Leonardo AB Tôrres, Luis A Aguirre, and Dennis S Bernstein. On unscented kalman filtering with state interval constraints. *Journal of Process Control*, 20(1):45–57, 2010.

[Timmons and Williams, 2020] Eric Timmons and Brian C. Williams. Best-first enumeration based on bounding conflicts, and its application to large-scale hybrid estimation. *Journal of Artificial Intelligence Research*, 67:1–34, 2020.

[Trave-Massuyes and Pons, 1997] Louise Trave-Massuyes and Renaud Pons. Causal ordering for multiple mode systems. In *Proceedings of the Eleventh International Workshop on Qualitative Reasoning*, pages 203–214, 1997.

[Tugnait, 1982] Jitendra K Tugnait. Detection and estimation for abruptly changing systems. *Automatica*, 18(5):607–615, 1982.

[Williams and Ragno, 2007] Brian C Williams and Robert J Ragno. Conflict-directed A* and its role in model-based embedded systems. *Discrete Applied Mathematics*, 155(12):1562–1595, 2007.