

# Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)

Ameet Kumar Jain<sup>a)</sup>

*Department of Computer Science, Johns Hopkins University, Baltimore, Maryland 21218*

Yu Zhou

*Department of Mechanical Engineering, State University of New York at Stony Brook, Stony Brook, New York 11794*

Tabish Mustufa

*Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218*

E. Clif Burdette

*Acoustic MedSystems, Inc., 2110 Clearlake Boulevard, Champaign, Illinois 61822*

Gregory S. Chirikjian

*Department of Mechanical Engineering, Johns Hopkins University, Baltimore, Maryland 21218*

Gabor Fichtinger<sup>b)</sup>

*Department of Computer Science, Mechanical Engineering, and Radiology, Johns Hopkins University, Baltimore, Maryland 21218*

(Received 4 January 2005; revised 15 August 2005; accepted for publication 13 September 2005; published 26 October 2005)

Intraoperative dosimetric quality assurance in prostate brachytherapy critically depends on discerning the three-dimensional (3D) locations of implanted seeds. The ability to reconstruct the implanted seeds intraoperatively will allow us to make immediate provisions for dosimetric deviations from the optimal implant plan. A method for seed reconstruction from segmented C-arm fluoroscopy images is proposed. The 3D coordinates of the implanted seeds can be calculated upon resolving the correspondence of seeds in multiple x-ray images. We formalize seed-matching as a combinatorial optimization problem, which has salient features: (a) extensively studied solutions by the computer science community; (b) proof for the nonexistence of any polynomial time exact algorithm; and (c) a practical pseudo-polynomial algorithm that mostly runs in  $O(N^3)$  time using any number of images. We prove that two images are insufficient to correctly match the seeds, while a third image renders the matching problem to be of nonpolynomial complexity. We utilize the special structure of the problem and propose a pseudopolynomial time algorithm. Using three presegmented images, matching and reconstruction of brachytherapy seeds using the Hungarian algorithm achieved complete matching in simulation experiments; and 98.5% in phantom experiments. 3D reconstruction error for correctly matched seeds has a mean of 0.63 mm, and 0.9 mm for incorrectly matched seeds. The maximum seed reconstruction error in each implant was typically around 1.32 mm. Both on synthetic data and in phantom experiments, matching rate and reconstruction error achieved using presegmented images was found to be sufficient for prostate brachytherapy. The algorithm is extendable to deal with arbitrary number of images without any loss in speed or accuracy. The algorithm is sufficiently generic to provide a practical solution to any correspondence problem, across different imaging modalities and features. © 2005 American Association of Physicists in Medicine. [DOI: 10.1118/1.2104087]

Key words: C-arm, fluoroscopy, seed matching, reconstruction, prostate brachytherapy, radiation planning

## I. MOTIVATION AND BACKGROUND

With an approximate annual incidence of 220,000 new cases and 33,000 deaths prostate cancer continues to be the most common cancer in men in the United States.<sup>1</sup> For several decades, the definitive treatment for low risk prostate cancer was radical prostatectomy or external beam radiation therapy,<sup>2</sup> but low dose rate permanent seed brachytherapy (shortly brachytherapy thereafter in this document) today can achieve virtually equivalent outcomes.<sup>3,4</sup> The success of brachytherapy (i.e., maximizing its curative force while

minimizing its co-morbidity) chiefly depends on our ability to tailor the therapeutic dose to the patient's individual anatomy. In contemporary practice, however, implant planning is based on idealistic preplanned seed patterns that, as 15 years of clinical practice has clearly demonstrated, are not achievable in the actual human body. According to a comprehensive review by the American Brachytherapy Society,<sup>5</sup> *the preplanned technique used for permanent prostate brachytherapy has limitations that may be overcome by intraoperative planning*. At the same time, continues the re-

port, *the major current limitation of intraoperative planning is the inability to localize the seeds in relation to the prostate*. There are excellent algorithmic and computational tools available today to optimize a brachytherapy treatment plan intraoperatively, thereby allowing for an improved dose coverage. These methods, however, critically require that the exact three-dimensional (3D) locations of the implanted seeds are precisely known with respect to the patient's anatomy.

**Transrectal ultrasound (TRUS) imaging:** Prostate brachytherapy is almost exclusively performed under TRUS guidance. While TRUS provides adequate imaging of the soft tissue anatomy, it does not allow for robust localization of the implanted brachytherapy seeds. Various researchers have tried to segment the seeds from TRUS images by linking seeds with spacers,<sup>6</sup> using x rays to initialize segmentation,<sup>7</sup> using vibro-acoustography<sup>8</sup> or transurethral ultrasound<sup>9</sup> as a new imaging modality, or segmenting them directly.<sup>10</sup> But even when meticulously hand-segmented, up to 25% of the seeds may remain hidden in ultrasound.<sup>11</sup> This necessitates the use of some other imaging method in intraoperative seed localization.

**Fluoroscopy:** The published history of C-arm fluoroscopy in brachytherapy originates<sup>12</sup> when it was first used as a solo guidance modality. Shortly after TRUS emerged as a primary image guidance modality, fluoroscopy became a secondary tool for gross visual observation. Mobile C-arms are ubiquitous in contemporary prostate brachytherapy, with approximately 60%<sup>13</sup> of the practitioners using it for qualitative implant analysis in the operating room. It is considered as the *gold standard* for intraoperative visualization of brachytherapy seeds. While several groups have published protocols and clinical outcomes favorably supporting C-arm fluoroscopy for intraoperative dosimetric analysis,<sup>10,14–22</sup> this technique is yet to become a standard of care across hospitals.

The ability to reconstruct and register the implanted seeds (that are visible in fluoroscope) to soft tissue anatomy (that is visible in TRUS) intraoperatively, would allow us to make immediate provisions for dosimetric deviations from the optimal implant plan. At the same time, quantitative use of fluoroscopy for dosimetric analysis has been hampered by a plethora of unresolved technical problems. The five major obstacles we face toward intraoperative dosimetry are: (a) C-arm distortion correction and calibration; (b) C-arm pose tracking; (c) seed segmentation; (d) seed matching and reconstruction; and (e) registration of C-arm to TRUS images.

Significant efforts have been made toward computational fluoroscopy guidance in general surgery,<sup>23,24</sup> developing various tools for distortion correction and calibration. However, C-arms available in most hospitals do not have encoded rotational joints, so one never knows where the fluoro shots are coming from relative to one another. We have addressed this issue by designing a fluoroscope tracking (henceforth FTRAC) fiducial, which is a radiographic fiducial system creating a unique projection image from each direction.<sup>25</sup>

Various methods partially dealing with C-arm calibration in brachytherapy have also been proposed,<sup>26–28</sup> while some others have suggested that it is redundant.<sup>29</sup>

**C-arm to TRUS registration:** Attempts have been made to relate fluoroscopic images to soft tissue anatomy.<sup>10,17,30–34</sup> Nevertheless, further research is merited since they are susceptible to various kinds of errors. We address this issue by the use of the FTRAC fiducial. It is not only capable of tracking the C-arm, but also of registering the C-arm to TRUS by a predetermined placement.

**Seed matching and reconstruction:** We assume that the seeds are 3D points and that their image locations are known, i.e., we do not address automatic segmentation for which methods are available.<sup>15,35–38</sup> Three-dimensional coordinates of the implanted seeds can now be calculated from multiple x-ray images upon resolving the correspondence of seeds, which remains the focus of this paper. Formalization of the seed-matching problem results in a high complexity search space of the order  $10^{150}$  and  $10^{300}$ , from two and three fluoroscopic images, respectively. Hence previously proposed seed-matching approaches have predominantly been heuristic explorations of the search space, with no theoretical assurance on the accuracy of the answer.

The early attempts<sup>39–41</sup> toward seed matching used three *coplanar images* (coplanar images are those where the implant and the three x-ray sources are approximately in the same plane). The images were divided into variable width bands, formed by comparing coordinates along the rotation axis. Furthermore, in order to make the bands, it was assumed that the seeds are near the iso-center of the C-arm or at least have similar magnifications in all the images. These methods are prone to calibration errors and become ineffective as the number of seeds increases. These ideas were further extended by accommodating for patient motion,<sup>42</sup> and yet all the seeds could not be reliably reconstructed. Further geometrical constraints were imposed by assuming that some of the seeds are in a straight line<sup>43</sup> or on quadratic curves,<sup>44</sup> which due to seed migration seems to be an assumption not supported by compelling evidence.

The first step toward mathematical formalization came with the construction of a cost matrix,<sup>45</sup> where exhaustive matching gave the lowest cost solution. Though it eliminated extraneous assumptions, it required impractical computational resources. A greedy randomized algorithm,<sup>46</sup> tested with various cost-metrics, was suggested to reduce the run-time. This method gives a different output for each run and is typically iterated a few hundred times, choosing the sequentially lowest cost. Though this method might provide an answer close to the correct match, its randomization does make any claim on the number of iterations required or proximity of the final answer. Fast-CARS<sup>47</sup> is another variant, which significantly improved the computational complexity, where “for each cost matrix an exhaustive match can be performed to obtain the best possible matches.” It reduced the run-time from  $O((N!)^2)$  to  $O((A!)^2)$ , where  $A$  is the average number of seeds in the band. Though it made the search faster, it still ran in exponential time. For example, if  $A=10$  then the num-

ber of computations would still be high at  $O(10^{14})$ .

Independently, a set of heuristic rules<sup>16</sup> were suggested for seed matching that attempted to reduce misclassifications. Simulated annealing<sup>36</sup> was proposed as an alternate technique to reach the global minimum. Another technique<sup>28</sup> was proposed that optimizes on seed positions and camera parameters, by generating simulated images and iterating them until they match the observed images. These optimization methods are prone to fall into local minima and may not be able to recover from them. A statistical simulation of seed reconstruction uncertainty was conducted,<sup>48</sup> but did not address the problem of seed matching. For completeness, we also mention that CT and MRI based techniques<sup>49,50</sup> are also proposed, but cannot be used intraoperatively and have poor resolution in the axial direction.

The matching problem is also prevalent in the computer vision community, where two-dimensional points are tracked and reconstructed to compute motion. Researchers have tried using noniterative greedy algorithms,<sup>51</sup> also incorporating spurious and hidden points.<sup>52,53</sup> Occlusion itself has also been a known problem.<sup>54</sup> These algorithms were optimized for a dense set of moving points, while specialized algorithms are used for sparse matchings,<sup>55,56</sup> which can also be used in pattern recognition<sup>57</sup> across images. These algorithms are usually catered to achieve real-time performance, as compared to a complete matching, and hence do not appear to be appropriate in a medical application.

The primary contribution of this work is a new theoretical framework for seed matching. The framework tackles issues of optimality and presents a practical algorithm that can be used. Moreover, the framework ensures a polynomial runtime of  $O(N^3)$  on the algorithm, an improvement over previous methods. While it has been motivated by prostate brachytherapy, it is a general purpose correspondence algorithm that appears to be usable in many synergistic problems. Besides the aforementioned problem, we are exploring matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL) for the matching between 3D clouds of objects, fiducial-based registration in CT/MRI imaging, and calibration of tracked ultrasound. In these applications, MARSHAL appears to be a potent tool that can replace traditional iterative closest point and other gradient descent optimizations whose exit criteria are governed by custom-tuned thresholds. Section II describes the theoretical foundation and the proposed algorithm. Section III discusses the experiments and results. In Sec. IV, we discuss the shortcomings of the algorithm and the future work. Finally, we conclude in Sec. V.

## II. MATERIALS AND METHODS

We convert seed-matching to network-flow-based combinatorial optimization. Our formulation has many salient features: (a) a global optimization using all seeds, as compared to a local seed clustering based approach; (b) exact solutions studied extensively by the computer science community; (c) addressing theoretically the achievable bounds by any algorithm; (d) guaranteed existence of a polynomial time solution




FIG. 1. The seed matching problem can be converted to a network flow graph. The best possible matching reduces to finding out the maximum flow with minimum cost.

for achieving the global minima for seed-matching from two images; (e) proof of the nonexistence of a polynomial time solution in case of more than two images; (f) a practical pseudo-polynomial algorithm using three images and mostly having an  $O(N^3)$  claim on the space-time complexity; (g) an algorithm that does not become slower as the number of images is increased. A polynomial-time algorithm is one that runs in a number of computations, polynomial in the size of the input. Further details are available in the literature.<sup>58</sup> We do not assume any prior information (e.g., inserted seed positions), the value of which is questionable due to intraoperative seed migration and tissue deformation.

### A. A generic network-flow-based formulation

A network flow formulation is created, where any flow in the network would represent a matching and the desired solution is the flow with minimum cost. Let  $N$  seeds be inserted and C-arm images  $I_1, I_2$  be acquired. Let  $s_{ij}$  be the position of the  $i$ th seed in  $j$ th image. We construct a directed network as shown in Fig. 1.

Sets A and B, each with  $N$  nodes, represent the two images  $I_1$  and  $I_2$ . While there are no edges within the set, directed edges (links) run from all vertices in set A to all vertices in set B. There are  $N$  links at source S, each link connecting to a node in A. Similarly each node in B is connected to sink T. The flow originates at S and ends at T, with each link allowing a flow of value 1 or 0, where 1 means that the edge is selected and 0 means that it is not. The problem is to efficiently compute a flow in the network that can achieve a total flow of value  $N$ .

It can be proved that any solution to the seed-matching problem is a solution to the flow problem and *vice-versa*. To have a net flow of  $N$ , each link connecting either the source or the sink has to support a flow 1. Now by the *conservation of flow* at each node, every node in set A will have to dispatch a unit flow to some node in set B. Moreover, each node in set B can accept only a unit flow, because any extra flow cannot be passed on to T and any deficiency would mean that T does not have a total flow of  $N$  units. The set of all links with nonzero flow provide a feasible matching. It can be verified that any matching of the seeds also provides a feasible flow. This proves that the flow problem is equivalent to the seed matching problem.

Simple combinations compute  $N!$  feasible solutions to the seed-matching problem, giving rise to  $N!$  feasible flows. To achieve the optimal solution, the link connecting seed  $s_{i1}$  to seed  $s_{j2}$  is assigned a cost  $C_{ij}$ . The cost  $C_{ij}$  represents the likelihood of seed  $s_{i1}$  matching seed  $s_{j2}$ , with the cost being 0 if they match perfectly and  $\infty$  (infinity) if they do not match at all. A popular example for the cost-metric is seed reconstruction accuracy. Further details are provided in Sec. II E. Any feasible flow has a net cost associated with it, the value of which is  $\sum_{i=1}^N \sum_{j=1}^N C_{ij} f_{ij}$ , where  $f_{ij}$  is the flow in link  $ij$  and  $C_{ij}$  is the cost of sending a unit flow along that link. Thus the seed-matching problem is reduced to finding the flow with minimum cost, and can be written down as

$$\min \sum_{i=1}^N \sum_{j=1}^N C_{ij} f_{ij},$$

$$\text{where } f_{ij} \in \{0, 1\}; \quad (1)$$

$$\sum_{j=1}^N f_{ij} = 1 \quad \forall i \quad \text{and} \quad \sum_{i=1}^N f_{ij} = 1 \quad \forall j.$$

The min-cost flow can be computed using cycle-canceling, successive shortest path, primal-dual, out-of-kilter or relaxation algorithms. These are all similar in their working and hence reducible to each other. Though they are straightforward to use, they run in pseudopolynomial time. The first weakly polynomial time algorithm was derived using an idea called scaling.<sup>59</sup> Capacity scaling, cost scaling, and double scaling algorithms that were developed on this idea are all weakly polynomial time algorithms. Repeated/enhanced capacity scaling and minimum mean-cost cycle algorithms achieve a strongly polynomial run-time. Alternately, linear-programming based techniques like simplex, interior point method or network simplex can also be used. Today, fast algorithms are available both freely<sup>60</sup> and commercially.<sup>61</sup> A comprehensive review of the above is available in the literature.<sup>62</sup>

We implemented the cycle canceling algorithm, which is the most straightforward of all the above. A feasible preflow is input as the current network flow. Then we generate the residual network, which is based on the difference between the network capacity and the current network flow, and search for negative cost cycles in the residual network. If there exists a negative cost cycle in the residual network, we decrease the cost by adding a flow along the negative cycle. Then we generate a new residual network based on this new flow in the network. We repeat this procedure until there is no negative cost cycle in the residual network. To find a negative cycle, we use the Bellman-Ford algorithm,<sup>58</sup> which runs in  $O(\text{Edges} * \text{Vertices})$  time. Thus the run-time for the cycle-canceling algorithm is  $O(\text{cost of the initial flow}) \times O(\text{Edges} * \text{Vertices})$ , which for integer costs will be pseudopolynomial. One way to boost performance is to use the negative cycle with minimum weight instead of any negative cycle, but that computation itself is NP hard. So researchers have suggested choosing the minimum mean-

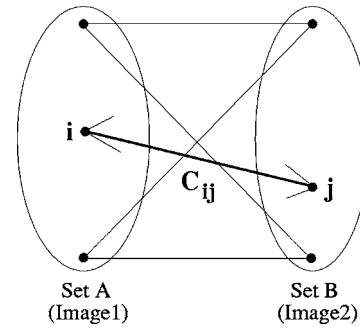


FIG. 2. Two image seed matching, modeled earlier as a min-cost max-flow problem, also reduces to the assignment problem. The best possible solution reduces to evaluating the minimum cost bipartite matching, which can be achieved in  $O(N^3)$  time using the Hungarian algorithm.

cost cycle or canceling several node-disjoint cycles at once, both of which have strongly polynomial run-times.

Hence the seed-matching algorithm will have a polynomial run-time. Though strongly polynomial time algorithms are available, in practice the basic cycle-canceling algorithm itself was sufficiently fast. This is due to the favorable problem structure arising from a bipartite network and epipolar constraints of x-ray imaging. The worst case theoretical run-time above is for generic networks. It should be noted that a significantly faster  $O(N^3)$  solution is possible if (a) all the nodes are available, (b) the number of nodes in both sets are equal, and (c) only unit flows are allowed. This is achieved by converting the generic network flow problem into a weighted bipartite matching problem, explained in the next section. Though we have implemented the cycle canceling algorithm, we have also tested and implemented a faster solution. However, this faster solution is not sufficiently generic to address the hidden seeds and other problems (Sec. IV).

## B. Seed matching from two images

When all seed locations in the two images are known, the minimum-cost maximum-flow formulation reduces further to the specific problem of minimum-weight matching in bipartite graphs, also known as the assignment problem (illustrated in Fig. 2) and can be solved very quickly. The problem is to find a minimum weight subset of edges such that all the vertices are covered exactly once. A real-world problem could be the assignments of jobs to workers in a factory. Each worker (set A) can do some jobs (set B) at a certain cost. We seek an assignment of jobs, such that all jobs are completed, each worker gets exactly one job, and the total cost is minimized. The assignment problem is also formalized by Eq. (1).

The assignment problem is solved in  $O(N^3)$  run-time by using the Hungarian algorithm.<sup>62-64</sup> Since the algorithm has been known for over four decades in the literature, we do not provide the full working of the algorithm, but instead just an outline. The  $N \times N$  cost matrix  $C$  is constructed. The final aim is to choose exactly one element from each row (and column) such that the sum of the elements has the lowest attain-



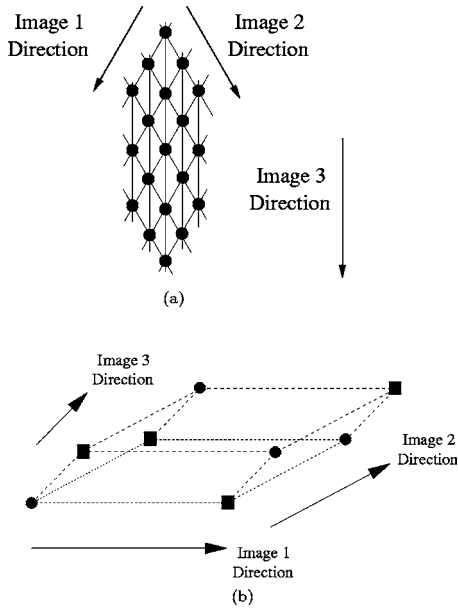


FIG. 3. (a) Three coplanar images also have seed constellation singularities. Multiple seed constellations will produce the same x-ray images. (b) Seed constellation singularity arising when three noncoplanar images are used. Two sets of reconstructions are viable for the same x-ray images. In fact this cube can be replicated to construct singularities that arise when more than three images are used.

able value. Thus, an equivalent matrix having at least one zero in each row (and column) is obtained by subsequent subtractions using the smallest element in each row (and column). This matrix is used to find a selection of zeros such that each row and column has exactly one zero. If this exists, then it provides the min-cost matching. If this does not exist, a *line covering* procedure is used to make an adjustment to the matrix and generate zeros in useful locations. This is iterated a maximum of  $O(N)$  times (proof available in the literature) until a solution results. The locations of all zeros provide the minimum weight matching. Thus the Hungarian algorithm provides the matching with the lowest possible cost.

**C. Theoretical foundations for uniqueness in matching**

It is well known that due to singularities in projective imaging, reconstructing seeds from two images is inherently inaccurate. A third image is required to remove this singularity. Further investigations reveal singular constellations even with three images. Previously proposed algorithms use three coplanar images, i.e., where the three x-ray sources and the implant are all in the same plane. Figure 3(a) illustrates the arising singular constellations for this choice of images. Due to the small size of the prostate, the x rays entering it are nearly parallel, the average variation being around  $2^\circ$ . The symmetry establishes that there are multiple constellations producing the same x-ray images, with the number of alternate constellations increasing with the number of seeds. Errors are further amplified with inaccuracies in segmentation, C-arm calibration and tracking. Though we used symmetric




FIG. 4. The three-image seed-matching problem reduces to the min-cost tripartite matching problem. Any seed in a given set is connected to every seed in the other two sets. The problem is NP-hard and no polynomial time solution is possible to compute the min-cost matching.

images, singularities can be shown in any three nonsymmetric coplanar images using the generic construction provided in the below.

Singularities though reduced, are not completely removed with noncoplanar images, as illustrated in Fig. 3. Since the x rays entering the prostate are nearly parallel, the three imaging directions define a parallelepiped, a 3D parallelogram. Four seeds sitting at the vertices of the parallelepiped can create a singularity. The likelihood of this happening increases for large implants, especially since the seeds are inserted in parallel straight lines. Thus a fourth image is advisable for large implants. Some singular constellations cannot be resolved by four images either. This constellation can be achieved by putting together two parallelepipeds along the fourth imaging direction, which will result in two singular constellations with eight seeds each. This process can construct singularities with any number of images, if taken arbitrarily. *Theoretically* seven images can resolve all singularities for purely convex objects. Similar results can also be derived when other properties of the object are known.<sup>65</sup> In brachytherapy however, three to four noncoplanar images should *practically* be sufficient.

The above implies that a robust algorithmic framework for seed matching using three or more images is essential. A modification of our proposed framework reduces three-image matching to the tripartite matching problem,<sup>64</sup> as illustrated in Fig. 4 and Eq. (2) ( $C_{ijk}$  is the  $N \times N \times N$  cost-matrix). Tripartite matching is similar to bipartite matching, except that it matches three sets instead of two. In addition, it assures a matching that is optimal, i.e., has a minimum cost. Multiple-image based seed matching similarly reduces to the multipartite matching problem. Though the bipartite problem is solvable in  $O(N^3)$ , the tripartite and multipartite problems are NP-complete. Moreover, finding the minimum weight tripartite matching is NP-hard. NP-complete are a class of dif-

difficult problems thought to be of nonpolynomial complexity. NP-hard problems are believed to be harder than NP-complete, where unlike NP-complete, even the validity of a given solution is not verifiable in polynomial time,<sup>58</sup>

$$\min \sum_{i=1}^N \sum_{j=1}^N \sum_{k=1}^N C_{ijk} f_{ijk},$$

where  $f_{ijk} \in \{0,1\}$ ;  $\sum_{i=1}^N \sum_{j=1}^N f_{ijk} = 1 \forall k$ ; (2)

$$\sum_{j=1}^N \sum_{k=1}^N f_{ijk} = 1 \forall i \quad \text{and} \quad \sum_{k=1}^N \sum_{i=1}^N f_{ijk} = 1 \forall j$$

Using another reduction, three-image seed matching reduces to the min-weight *clique* problem<sup>58</sup> (a clique is a subgraph with edges between all the vertices). If the images are taken from arbitrary directions, then the orientation of epipolar lines is unconstrained. Since the implant has  $N$  densely packed seeds roughly shaped as an ellipsoid, any epipolar line connecting the projection of a seed to its x-ray source is expected to come in the vicinity of  $O(N^{1/3})$  seeds in the 3D implant. Note that this estimate is an expected bound and could vary significantly across individual seeds. When two images are used, all seeds in the vicinity of the plane connecting the seed and the two x-ray source locations would become feasible. Thus, with unconstrained image orientations, each seed in an image *statistically* has an expected  $O(N^{2/3})$  feasible correspondences in the other image. This implies that two-image matching has an expected  $O(N^{5/3})$  feasible doublets. Similarly, it can be established that three-image matching has an expected  $O(N^{4/3})$  feasible triplets.

This reduction has also been observed experimentally where two images resulted in about 2000 feasible doublets for a 100 seed phantom, while a third image reduced it to about 400 triplets. As the number of images is increased, the number of feasible tuples will converge close to  $O(N)$ , implying a decrease in the number of alternate implants. This number can potentially reduce to exactly  $N$ , when many well-chosen images are used. It should be observed that even having only  $(1 + \epsilon) \times N$  feasible tuples results in an exponential search space of  $O(2^N \times (\epsilon - 1)^N)$  for large  $\epsilon$  and  $O([2/(\epsilon - 1)]^{\epsilon - 1N})$  for small  $\epsilon$ . A formal proof for the non-existence of a polynomial-time algorithm can be shown by constructing a graph with all feasible tuples as vertices, and compatibility between triplets as edges. The solution will be equivalent to finding a min-weight clique of size  $N$ . Finding a feasible clique in a general graph is itself NP-complete, while finding the clique with minimum cost<sup>58</sup> is NP-hard.

### D. A practical algorithm for three-image seed matching

Though still an open problem in complexity theory, it is widely believed that no polynomial time algorithm can solve NP-complete or NP-hard problems.<sup>58</sup> The field of *approximation algorithms* could design quick algorithms that have

theoretical bounds on the error for both run-time and the final solution. Unfortunately, these algorithms typically work on generalized graphs, and may not necessarily incorporate the inherent structure of our problem. Hence we design our own algorithm that incorporates the physics of projective imaging. We propose a practical solution for matching and reconstruction of brachytherapy seeds using the Hungarian algorithm, abbreviated henceforth in the paper as MARSHAL.

The intuition behind MARSHAL can be summarized into: (a) the original tripartite matching can be projected into *inspired* bipartite matchings; (b)  $s_{i1}$  matches  $s_{j2}$  only if they have a counter-part in the third image; (c) though a low  $C_{ij}$  did not force  $C_{ijk}$  to be low for some  $k$ , a *low*  $C_{ijk}$  does force  $C_{ij}$ ,  $C_{jk}$ , and  $C_{ki}$  to be *low*. It is very common in the two-image case for the epipolar line of  $s_{i1}$  to pass through  $s_{j2}$ , even when  $s_{i1}$  and  $s_{j2}$  are severely mismatched. Thus in a matching between images 1 and 2,  $C_{ij}$  should be low if and only if there exists a  $k$  such that  $C_{ijk}$  is low. If  $C_{ijk}$  is high for all  $k$ , then  $C_{ij}$  should be given a high value. This assignment of  $C_{ij}$  is different from the purely two-image case, where  $C_{ij}$  could be low even for a strong mismatch. Thus  $C_{ij}$  needs to incorporate information from the third image to remove inherent two-image singularities. Since the solution of bipartite matching should be as close as possible to that of the tripartite matching, the  $L_{-\infty}$  projection (minimum/best value) is used.

$$C_{ij} = \min\{C_{ijk}/k = 1, \dots, N\}. \quad (3)$$

MARSHAL *projects* the original tripartite problem into three distinct bipartite problems by the appropriate projection of the costs. This is similar to projecting the *minimum value* along the rows/columns of the three-dimensional cost matrix  $C_{ijk}$  to obtain three two-dimensional cost matrices. The theoretical framework of bipartite matchings allows the computation of the best possible matching for each bipartite problem in  $O(N^3)$  run-time. The solutions of the bipartite matchings are then *integrated* to obtain a solution to the original tripartite problem. Hence MARSHAL acts as a bridge between a purely theoretical framework and an inherent structure in the problem.

The flowchart for MARSHAL is shown in Fig. 5. Using Eq. (3), the respective cost matrices for the three image pairs are obtained. These are used to achieve three independent bipartite matchings  $M_1\langle i, j, - \rangle$ ,  $M_2\langle -, j, k \rangle$ , and  $M_3\langle i, -, k \rangle$ . Loops are created next, i.e., if  $\langle i1, j1, - \rangle, \langle -, j1, k1 \rangle, \langle i2, -, k1 \rangle, \langle i2, j2, - \rangle, \langle i1, -, km \rangle$  are matchings, then  $\langle i1, j1, k1, i2, j2, k2, \dots, km, i1 \rangle$  is a loop of size  $m$ , each loop ending on the seed it started from. For example, if for the first seed the matchings are  $M_1\langle 1, 1, - \rangle$ ,  $M_2\langle -, 1, 1 \rangle$ ,  $M_3\langle 1, -, 1 \rangle$ , then  $\langle 1, 1, 1 \rangle$  is a loop of size 1 representing *seed1*. The majority of the loops will have a size of 1, which can be declared matched. This claim is also experimentally validated in Sec. III C.

Sometimes matchings get flipped, resulting in loops of size greater than 1. For example, if  $s_{i3}$  had instead matched  $s_{j1}$  in the above-noted example. i.e., if the matchings

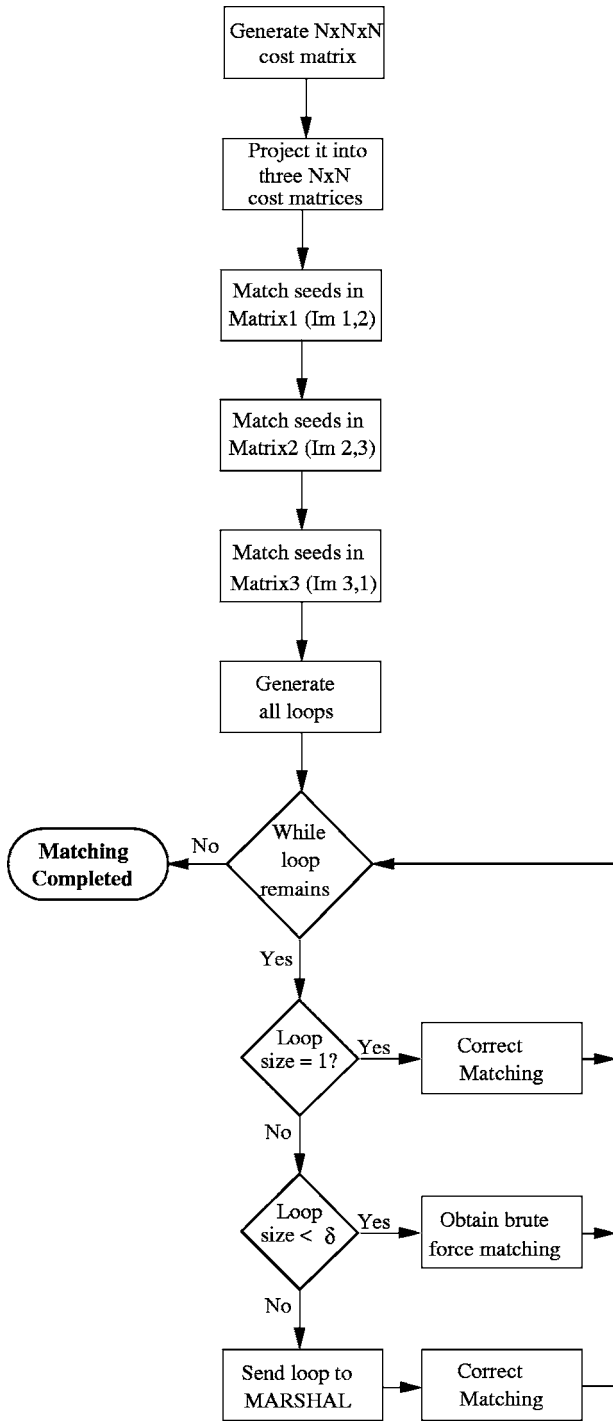


FIG. 5. A flowchart explaining the working of MARSHAL. It runs in  $O(N^3)$  time for good data sets.

were  $M_1\langle 1,1,-\rangle$ ,  $M_2\langle -,1,1\rangle$ ,  $M_3\langle 2,-,1\rangle$ ,  $M_1\langle 2,2,-\rangle$ ,  $M_2\langle -,2,2\rangle$ ,  $M_3\langle 1,-,2\rangle$ , then  $\langle 1,1,1,2,2,2,1\rangle$  would be a loop of size 2 representing seed1 and 2. Since each loop by definition is a self-contained subset of seeds with no relation to other loops, they only need to be matched among themselves to obtain the correct final answer. These are typically small loops, the optimal match for which can easily be obtained by a brute-force search with  $O(m! \times m!)$  run-time.

Theoretically, for error-prone data the size of the loops could grow large, though practically we never saw it happen. Nevertheless, brute force search is done only for loops of size  $m < \delta$ , which is a predetermined threshold based on the largest  $m$  that achieves a very fast ( $< 1$  s) run-time. Loops of size greater than  $\delta$  are recursively broken down using MARSHAL, which would run three bipartite matchings within this subset and return the correct match. In effect, all large loops will be broken down into smaller and smaller loops, until everything is matched. Thus MARSHAL achieves the correct seed matching and still practically runs in  $O(N^3)$ , a significant improvement when compared to previous methods.

It should be noted that there can be another way to compute loops in the above algorithm. If  $\langle i,j,-\rangle$  match after the first bipartite matching, then their counterpart  $k$  in  $Image_3$  could also be chosen. We do not do this, since the first bipartite matching does not ensure that all  $k$ 's are chosen, to solve which the full tripartite problem would have to be solved. Moreover,  $\langle i_1,j_1,-\rangle$  and  $\langle i_2,j_2,-\rangle$  can map on the same  $k$ , or some  $k$  can be left untouched after the first bipartite matching. Thus the first bipartite matching attempts to compute as accurate as possible a correspondence in the first two images, while the next two matchings will later solve for any inconsistencies.

Note that the computation in Eq. (3) takes  $O(N^3)$  time. We modify it to run in  $O(N^2)$  time. To compute  $C_{ij}$ ,  $s_{i1}$  and  $s_{j2}$  are used to reconstruct the 3D point  $S'_{ij}$ , which is then projected on  $Image_3$ . Let  $k$  be the closest point in  $Image_3$  to this projection. Now  $S_{ijk}$  is computed and its projection error (PE) calculated, which is assigned to  $C_{ij}$ . Though this  $C_{ij}$  does not strongly satisfy Eq. (3), it weakly satisfies it. Any consistent  $\langle i,j \rangle$  will choose the correct  $k$ , while an inconsistent  $\langle i,j \rangle$  does not have any correct counter-part  $k$ , and will anyway be eliminated in the bipartite matching. Thus our  $C_{ij}$  exactly satisfies Eq. (3) in relevant cases, and will become irrelevant otherwise.

### E. Cost-metric

A comparison proves that some metrics are more effective than others.<sup>46</sup> The metric should incorporate all the available information, making the global minimum sharper and the algorithm robust. Hence  $C_{ijk}$  should ideally represent the inconsistency among seeds  $s_{i1}$ ,  $s_{j2}$ , and  $s_{k3}$ , with 0 indicating least inconsistent. The basis for most popular choices of  $C_{ijk}$  are reconstruction accuracy (RA) and projection error (PE). We do not formally compare the various metrics that we have tried, but instead provide the intuition behind the preferred one.

RA, used by most researchers, is computed by first calculating the equation of the three lines that join each projection to its respective x-ray source. Due to various errors these lines never intersect, creating a need for a symbolic intersection. As explained in Sec. II F, a closed form solution for RA that minimizes the  $L_2$  norm of the distance vector is used in this publication. Though they do not have an easy closed form solution,  $L_1$ ,  $L_3$ ,  $L_\infty$ , or other norms can also

be minimized ( $\|v\|_{L_1} = \sqrt{\sum_k v_k^2}$ ). The  $L_\infty$  norm would be the most robust, since the higher the order of the norm, the more sensitive it is to the smallest of inconsistencies. Thus, it measures the accuracy with which the 3D point can be reconstructed. PE, an alternate metric, can be computed by extending RA. The 3D reconstructed point can be projected back into each image, and the mean distance between the projected location and the observed location of the seed constitute PE. The symbolic intersection of the 3D point can be computed by minimizing any choice of norm ( $L_1-L_\infty$ ), and could possibly affect a change in the performance. We have minimized the  $L_2$  norm to compute the 3D intersection, and used its PE as the metric.

In our experience, PE fares significantly superior as a cost-metric than RA. Though PE for each seed is its RA magnified, different seeds will have different magnifications (depending on depth). If two seeds have similar RA but differing PE, the seed with lower PE should ideally be chosen by the optimization. On the other hand, an algorithm minimizing RA as a metric will not necessarily choose a correspondence that minimizes the deviation of reconstruction from observation. Thus PE is a stronger reflection of the observed data and leads to better convergence. In general, any cost-metric that directly measures the deviation from the observation performs superior to a metric that does not. A significant improvement in matching was observed by switching from RA to PE. A combination of the above-mentioned ideas is to choose the intersection so as to minimize the  $L_\infty$  norm of PE. Thus the most sensitive norm ( $L_\infty$ ) is minimized on the observations (seed locations on image), to compute the 3D intersection. Unfortunately, it requires a nonlinear optimization due to a lack of a closed form solution. Though we have not yet used it, this metric is expected to give the best performance.

**F. Seed reconstruction**

To compute  $\mathcal{C}$ , we need to compute the 3D intersection of the corresponding straight lines in space. Due to various errors these straight lines never intersect, forcing us to compute a symbolic 3D intersection point. The symbolic intersection is typically defined as the global minimum of an error function. Here we propose a simple and quick method that minimizes the  $L_2$  norm of Euclidian distance from the intersection point to the lines.

Let the total number of 3D straight lines be  $m$ , with line  $i$  defined as having unit direction  $l_i (a_i, b_i, c_i)$  and a point  $p_i$  on it, as shown in Fig. 6. Let  $P (x, y, z)$  be the representative intersection of these  $m$  lines. Let  $d_i$  be the Euclidian distance of  $P$  from line  $i$ . Thus by definition,  $P$  achieves the minimum  $L_2$  norm for the vector  $(d_1, d_2, \dots, d_m)$ . In other words, we need to find a  $P$  such that it minimizes a function  $\mathcal{F}$ . Now, it can be easily computed that




Fig. 6. Three matching points, in general, do not intersect due to various errors in segmentation, C-arm pose, and calibration. A symbolic intersection is calculated by finding the point with minimum sum of square distances from the lines.

$$\begin{aligned} \mathcal{F} &= m \times RA^2 \\ &= \sum_{i=1}^m d_i^2 \\ &= \sum_{i=1}^m \|(P - p_i) \times l_i\|^2 \\ &= \sum_{i=1}^m (P - p_i)^T \begin{bmatrix} b_i^2 + c_i^2 & -a_i b_i & -a_i c_i \\ -a_i b_i & a_i^2 + c_i^2 & -b_i c_i \\ -a_i c_i & -b_i c_i & a_i^2 + b_i^2 \end{bmatrix} (P - p_i) \\ &= \sum_{i=1}^m (P - p_i)^T A_i (P - p_i), \end{aligned} \tag{4}$$

where  $T$  represents the transpose. Now  $P$  minimizes  $\mathcal{F}$ . Hence we can write,  $\partial \mathcal{F} / \partial x = \partial \mathcal{F} / \partial y = \partial \mathcal{F} / \partial z = 0$ , which leads us to

$$P = \left[ \sum_{i=1}^m A_i \right]^{-1} \times \left[ \sum_{i=1}^m A_i p_i \right]. \tag{5}$$

This is the final seed coordinate. We can see that it can be computed very quickly by a few summations followed by a  $3 \times 3$  matrix inversion. It should be noted that  $P$  is chosen so as to minimize RA.

**G. Seed matching from four or more images**

In some rare cases of extremely large/dense implants, an additional fourth image may be desired. One approach is to match all possible image pairs and perform a combinatorial search on the loops that are formed. Though this is accurate, the run-time complexity is dependent on the number of images, making the algorithm slow. It was earlier observed that the number of feasible  $m$ -tuplets rapidly comes closer to  $N$ . This property can be used to extend MARSHAL. The algorithm chooses three images at random as primary and the rest  $(m-3)$  as secondary. Using the primary images, it runs exactly as described in Fig. 5, except that each bipartite problem is a projection of the  $m$ -partite problem. In other words  $C_{ij}$  incorporates information from all the  $m$  images, as shown as follows:






FIG. 7. Two seeds nearly overlap in the right image. Though both matchings A and B reconstruct the seeds in the correct location, only one of them is theoretically correct. We define the other as mismatched, so as to make the seed matching error analysis extremely rigorous.

$$C_{ij} = \min\{C_{i_1 i_2 \dots i_m} / i_1 = i; i_2 = j \text{ and } i_3, \dots, i_m \in \{1, 2, \dots, N\}\}. \quad (6)$$

Extended-MARSHAL will still run in  $O(N^3)$  time. In comparison to three-image matching, using four images will in fact have a faster run-time due to a decrease in loop-size and due to a better conditioned bipartite matching. An improvement in accuracy is also registered owing to the inherent structure in the additional information.

#### H. Guidelines for performance evaluation

Since MARSHAL is essentially a correspondence algorithm, the % of correctly matched seeds (mean, STD) is the optimal metric to evaluate its performance. However, it should be noted that due to the *overlap-problem*, the % of declared mismatches would always be an upper-bound. As illustrated in Fig. 7, both matchings A and B give the correct 3D seed coordinates in a *practical sense*, though only matching A is correct in a *theoretical sense*. Owing to our rigid error analysis criterion, we declare B to be an incorrect matching, in spite of it being correctly matched in a practical viewpoint. This will be evident from the low reconstruction errors for mismatched seeds, which should be comparable to that of correctly matched seeds.

Error analysis is sometimes also done using PE or RA. As theoretically described in Sec. II C, a mismatch might still result in a low PE/RA. This makes performance evaluation based on PE/RA inaccurate. If available, reconstruction error (RE), the 3D distance between the true and the reconstructed seed location, is an excellent tool for measuring performance. To measure any systematic bias in implant reconstruction or any change in shape, we use the relative reconstruction error (relative RE). It is the 3D error when the reconstructed implant is compared to a shifted (translated, rotated) version of the ground truth implant.

Typically RE is averaged for all the seeds, leading to potential information misrepresentation. For example, a 100-seed implant can have 90 seeds correctly matched with 0.25 mm RE, and 10 severely mismatched seeds with 5 mm RE. Though the *average* RE in this case is only about 0.7 mm, the 10 mismatched seeds could be hazardous. Thus a small average RE alone does not necessarily imply a good matching. A *separation of matched and mismatched seeds is necessary*. Note that the distinction between matched and mismatched seeds should never be established by using PE/

RA, but only by an accurate ground truth. RE/PE/RA in addition to the above, also reflect the accuracy of C-arm tracking and seed segmentation. In this light, MARSHAL is primarily evaluated using % matching and RE for mismatched seeds.

Beside establishing tools for performance evaluation, it is also important to find factors that could adversely affect a seed matching algorithm. This is crucial, since no polynomial algorithm can assure optimality for more than two images. The primary factors are the implant seed density and the size of the prostate. The higher the density or bigger the prostate, the greater the possibility of mismatches. Besides, most available C-arms are of moderate quality, with little reliability of the system supplied parameters. This includes encoder readings (pose) and parameters like image warp, focal length, and image origin. Another factor that could affect performance is the angle of separation between the images. Thus sensitivity to these factors should be considered when evaluating any new algorithm.

### III. RESULTS AND DISCUSSION

#### A. Two-image versus three-image matching

By converting the two-image matching problem to a bipartite matching problem, we have assured that the minimum-cost match from all the possible  $N!$  matches can be computed in  $O(N^3)$  run-time. This proves that there cannot exist a better match, implying that no other algorithm can produce a match with a lower cost. It should be noted that the best possible match is not representative of the correct match, which is due to the fact that two images are intrinsically insufficient and do not contain the requisite information. Furthermore, the minimum-cost match is dependent on the choice of cost-metric, i.e., a superior choice for the cost-metric will result in a superior match. Hence we have proved that given two images and a choice for the cost-metric, MARSHAL computes the best possible match (global optima) in polynomial time. Furthermore, we have shown that information from two images is inherently ill-posed and that the best match does not correspond to the correct match.

This is also validated using quantitative experiments. Synthetic C-arm images of clinically realistic brachytherapy implant scenarios were created, and the reconstructed and actual seed locations were compared. Implants with the number of seeds varying from 60 to 150 in a 50 cc prostate were tested. Based on our experience, the focal length was chosen to be 1000 mm, the pixel size was chosen to be 0.25 mm and the center of the image was chosen to coincide with the imaging origin. The region of use was at about a distance of 2/3rd the focal length from the x-ray source. The C-arm images were all contained inside a  $15^\circ$  cone around the AP axis. To reflect the nature of the best possible results, no noise was added in these simulations. Table I shows the performance on this synthetic data.

It can be noticed that two images are clearly insufficient, correctly matching only 85.2% of the seeds even in the absence of any noise. The inherent singularity in the two-image case is indicated by the low PE and high RE for the mis-

TABLE I. MARSHAL performance using synthetic data. Two images are clearly insufficient, which is clear from the low projection error yet high reconstruction error for the mismatched seeds. All seeds are matched when three images are used.

		Number of seeds									
		Two images					Three images				
		60	80	100	120	150	60	80	100	120	150
Matching rate (%)	Match	91.3	90.1	80.9	82.8	80.8	100	100	100	100	100
Reconstruction error (mm)		0.18	0.15	0.18	0.20	0.21	0.07	0.07	0.07	0.07	0.07
Reconstruction error (mm)	Mismatch	22.1	21.4	22.1	23.6	21.9	...	...	...	...	...
Projection error (mm)		0.03	0.04	0.04	0.04	0.04	...	...	...	...	...

matched seeds. Thus, when only two images are used, there exist alternate implants that produce the same x-ray images (low PE of 0.04 mm), but reconstruct the mismatched seeds at the wrong 3D location (high RE of 22.2 mm). Another way to look at the results is that the desired matching of the seeds has a cost associated with itself. This cost is close to 0 (in practice it is always a little more than 0 due to pixelization errors). Due to inherent singularities, a significant number of alternate implant constellations also have a cost very close to 0. Theoretically, all these are viable implants. Nevertheless, even a pixel size of 0.25 mm results in many of these alternate implants having a cost lower than the desired implant, which is then chosen by the algorithm. Thus, any algorithm which is minimizing the error from observed data will choose an incorrect matching. This is not the fault of the matching but rather a deficiency of the input. The suggested approach, therefore, would be to increase the number of images. Similar conclusions for two-image matching can also be drawn from phantom experiments, complete results of which are available in Sec. III C.

In comparison, using a third image matched all seeds correctly on noise-free data. Thus while two images are insufficient, three are practically adequate. As mentioned in Sec. III C, three-image matching reduces to the tripartite matching problem. Multiple-image based seed matching similarly reduces to the multipartite matching problem. Though the bipartite problem was solvable in  $O(N^3)$ , the min-cost tripartite and multipartite problems are NP-hard, meaning no polynomial algorithm will be able to prove that its solution is the best achievable. Thus this proves that there cannot exist any polynomial time algorithm which can solve the matching problem using three or more images. It should be noted that algorithms providing practically acceptable solutions (without any claim on absolute optimality) can still exist.

## B. Simulations

Studies were conducted on synthetic images to analyze the effects of various governing parameters. MATLAB software was created to model x-ray imaging. Given C-arm parameters and implant details, it generated synthetic images

and exact locations of seeds in the images. Random error was modeled using a uniform probability density function, i.e., a 1 mm error means that a maximum error of *magnitude* 1 mm was added, having a uniform probability distribution. The C-arm geometry was the same as in the previous section. Individual parameters were changed to understand the sensitivity of MARSHAL. It should be noted that only three images were used to evaluate the correspondences in all the simulation experiments. Using four or more images is expected to improve performance, though not evaluated in simulation experiments.

Since MARSHAL is essentially a correspondence algorithm, to evaluate performance on simulated data, we computed only the % of correctly matched seeds (mean, STD). RE for matched seeds is not plotted, as it is a function of the particular methods used for C-arm tracking and segmentation, which are not the focus of this paper.

**Segmentation:** Seed segmenting from the x-ray images is one of the most important sources of error. To study the sensitivity of MARSHAL to segmentation, segmentation error ranging from 0 to 2 mm in steps of 0.25 mm was added to each seed coordinate in the synthetic data sets. The computed correspondence was compared to the correct known correspondence, and the % of correctly matched seeds was evaluated. The seed density was varied from 1.5 to 1.9 seeds/cc in steps of 0.1, while the prostate volume was varied from 45 to 55 cc in steps of 5 cc. The averaged results (mean and STD) from a total of 54 000 iterations are plotted in Fig. 8(a).

The results indicate that even with segmentation errors as high as 1 mm, the mean and STD for % matching is better than 97% and 2%, respectively. With 1–2 pixel segmentation error, the average matching rate is practically 100% in all data sets. Segmentation error needs to increase beyond 1.5 mm for the mean matching rate to fall below the 95% mark. Moreover, it should be noted that the declared mismatches are predominantly due to the overlap problem, which becomes clear from the low RE for mismatched seeds. It has an average of 0.52 mm for all the data sets. The maximum is attained for data sets with 2 mm segmentation error

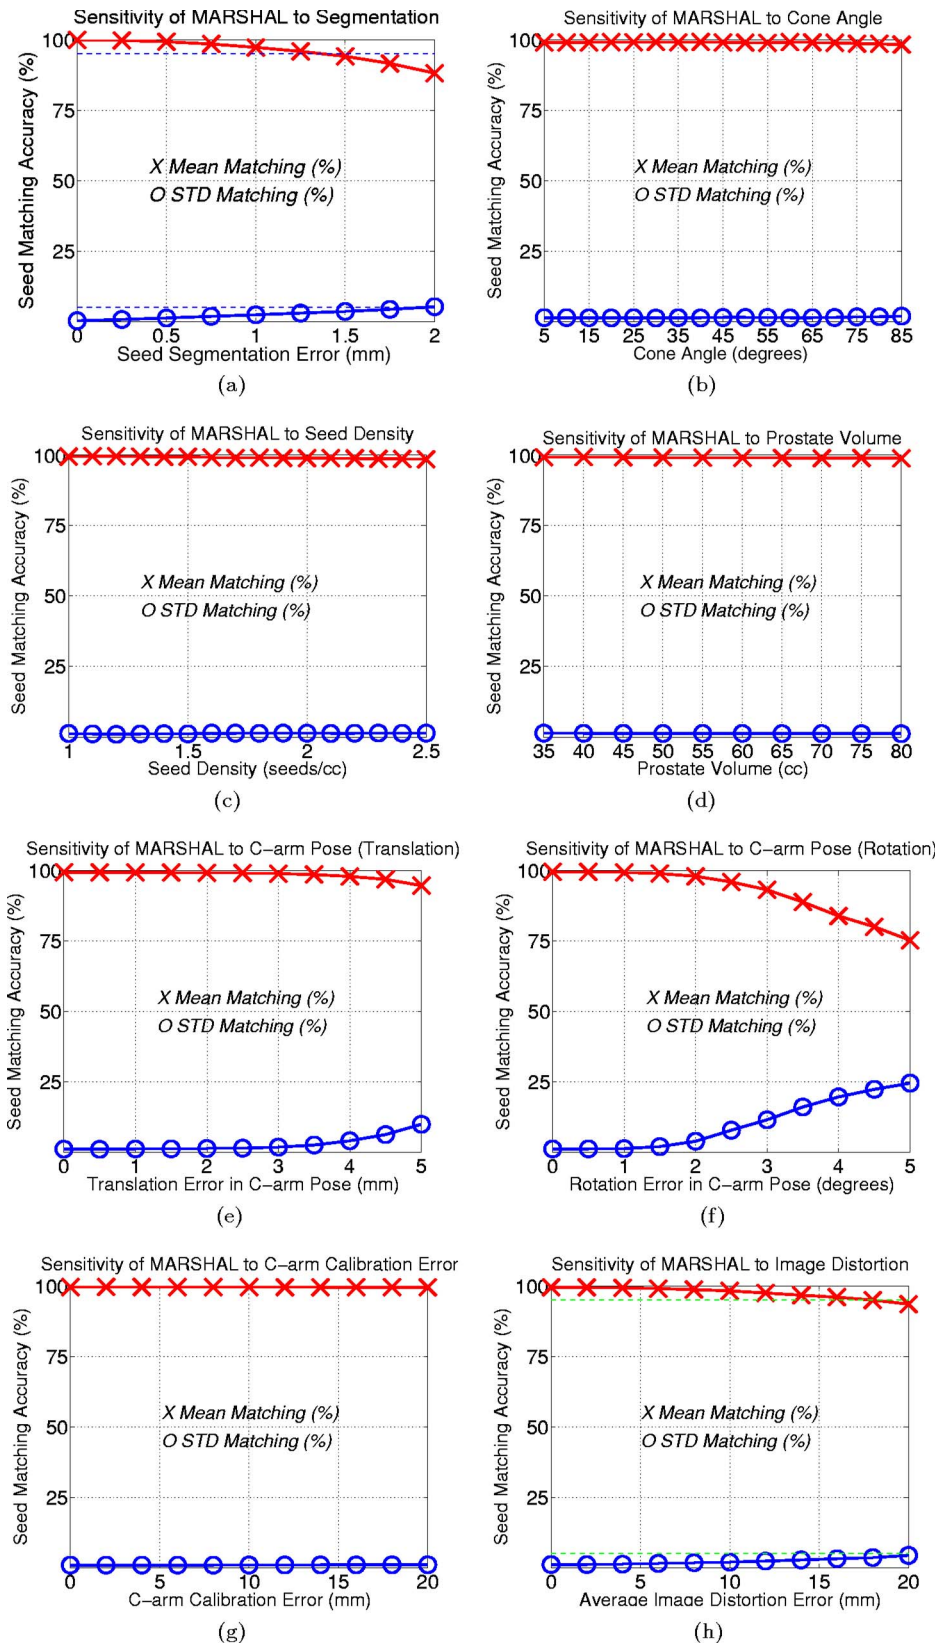


FIG. 8. Error in the performance of MARSHAL: (a) segmentation error; (b) cone angle variation; (c) seed density variation; (d) prostate volume variation; (e) C-arm pose translation error; (f) C-arm pose rotation error; (g) C-arm calibration error; and (h) C-arm distortion error.

at 1.51 mm, which is natural due to the presence of noise in segmentation, as compared to an error in matching. MARSHAL requires seed segmentation errors to be under 3–4 pixels (~1 mm) for optimal performance. For the re-

mainder of the simulations, we fixed a segmentation error of 0.5 mm (1–2 pixels) to be added to all the datasets.

**Seed density and prostate volume:** Depending on the radio-activity of the chosen seed, the seed density is com-



puted. From some preliminary calculations, seed density typically ranges from 1.5 to 2 seeds/cc.<sup>66</sup> An increase in seed density could make the matching process more difficult, since the seed projections in the x-ray image gets more and more congested. Thus it is extremely necessary to evaluate the sensitivity of any algorithm to seed density. In a similar manner, a larger prostate would need a greater amount of dose, resulting in a greater number of seeds. An increase in the number of seeds can potentially be of concern as the number of feasible matches would increase. Matching algorithms are expected to be less sensitive to prostate volume in comparison to seed density. Nevertheless, they should be evaluated in view of volume changes.

The seed density was changed from 1 to 2.8 seeds/cc in steps of 0.1 and the prostate volume was changed from 35 to 80 cc in steps of 5. As mentioned before, segmentation error was kept constant at 0.5 mm. A total of 23 750 data sets were evaluated. Averaged results for matching rate, as a function of seed density and prostate volume are plotted in Figs. 8(c) and 8(d), respectively. It can be noticed that both the mean and STD are stable throughout with a variation of less than 1%. The mean matching rate is 99% or above, with a few seeds in some cases being counted as mismatched due to the overlap problem. Thus MARSHAL can be declared robust to variations in seed density and prostate volume.

**Separation angle:** The three or more x-ray images of the implant are typically taken with a rotation motion of the C-arm. Some earlier work had made certain assumptions on the path of the C-arm, while some other work used orthogonal images to do the reconstruction. It is desirable to not have any constraints on how the images are taken. Moreover, in a clinical setting the C-arm would have only a limited mobility inside a 25° cone around the AP axis. Thus, it is necessary to validate an algorithm to the angular separation between the images. We tested the algorithm on various data sets with images taken on a cone around the AP axis. The implant was kept close to the center of the cone, which was in proximity of the iso-center. The cone angle was varied from 5° to 85° in steps of 5°. Six images were taken on each cone and all 20 three-image combinations used. Seed density was varied from 1.8 to 2 seeds/cc, prostate volume from 45 to 55 cc, and segmentation error was constant at 0.5 mm.

Averaged results from a total of 77 400 data sets are plotted in Fig. 8(b). Mean matching rate is fairly constant remaining always over 99%. There does not seem to be any variation in % matching due to a variation in the angular separation of the images. Thus it can be concluded that MARSHAL is robust to image-separation, and any three images in the available workspace should suffice. It should be noted however that, though MARSHAL can match the seeds correctly it does not have a bearing on RE, which for each seed will be a function of cone angle, C-arm tracking and seed segmentation. Thus a larger image separation is desired since it improves RE.

**C-arm pose tracking:** The accuracy of C-arm tracking has a definite effect on the ability to correctly match the seeds. We separately evaluate the performance of MARSHAL to errors in C-arm translation and rotation. Statisti-

cally uncorrelated pose errors were added to all three images. Errors in translation varied from 0 to 5 mm in steps of 0.5 mm, while in rotation varied from 0° to 5° in steps of 0.5°. It was observed in our fluoroscope tracking (FTRAC) fiducial that translation errors in depth are always significantly greater than those parallel to the plane<sup>25</sup> (by a factor of at least 5). This statistical bias in translation was incorporated in the generation of the datasets. There was no significant bias observed in the rotation errors.

Averaged results for translation and rotation from a total 44 000 data sets are shown in Figs. 8(e) and 8(f), respectively. It can be noted that translation errors up to 3 mm can still correctly match over 99% of the seeds. This quickly drops to 95% when the error reaches 5 mm. When compared to translation, MARSHAL is more sensitive to errors in rotation. Errors up to 1.5° in rotation correctly match over 99% on the seeds, with the performance dropping quickly when the error is greater than 3°. It can be concluded that MARSHAL is accurate and robust when the C-arm is tracked to an accuracy of 3 mm in translation and 2° in rotation. This is significantly larger than accuracies offered by the FTRAC (0.56 mm in translation and 0.33° in rotation)<sup>25</sup> or other fiducials.<sup>34</sup>

**C-arm calibration:** Quantitative fluoroscopic reconstruction requires accurate calibration of the C-arm imaging model parameters. Since the parameters vary from pose to pose, this is typically done for every pose from which the image is taken. The five intrinsic camera parameters are the pixel sizes (two parameters) and the 3D location of the x-ray source with respect to the image (three parameters). Since pixel sizes are fixed throughout the life of the C-arm, we evaluate MARSHAL only with respect to the focal spot. Independent and uncorrelated calibration errors of up to 20 mm were added to all three images, and the averaged % matching from 102 500 data sets are plotted in Fig. 8. It can be seen that the matching rate stays practically at about 100%. Thus MARSHAL is robust to C-arm calibration. These results triggered us to question the relevance of C-arm calibration in brachytherapy, leading to both theoretical and experimental bounds on the error.<sup>29</sup>

**C-arm image distortion:** Most C-arms exhibit a significant amount of nonlinear distortion in the images. Moreover, the amounts of distortion vary with pose, time, and location. Distortion being a global and nonlinear transform, consistently shifts the seed coordinates in the image, adding an error to any algorithm that relies on epipolar constraints. Thus distortion correction needs to be accommodated for in every image, which is a cumbersome process. If any matching algorithm is stable in the presence of distortion, then a one-time distortion correction procedure might also suffice. MARSHAL solves the matching problem globally, as compared to decreasing the complexity to smaller epipolar constraint based subsets, and is well suited to work on distorted data sets.

MARSHAL was tested for sensitivity to distortion. A randomly generated fifth degree radial distortion function was used to distort each image, and the matching computed on this data set. The implants were not confined to be in the






FIG. 9. (a) An image of the seed phantom attached to the FTRAC fiducial. The seed phantom can replicate any implant configuration, using the twelve 5 mm slabs each with over a hundred holes. (b) A typical x-ray image of the combination.

center of the image. Averaged % matching from 27 500 data sets are plotted against the mean image-distortion in Fig. 8. Even with 15 mm distortion errors, % matching is above 95%. Average image distortion of up to 5 mm matches over 99% of the seeds correctly. Thus MARSHAL seems to be fairly robust to moderate distortion errors. Again, it fares well since it solves the matching problem as a global combinatorial optimization. Similar results are also available from phantom experiments in Sec. III D.

### C. Phantom experiments

After the simulations, experiments were conducted on a precisely fabricated seed phantom, constructed using acetol. The FTRAC fiducial<sup>25</sup> was used to track the C-arm (accuracy of 0.56 mm translation and 0.33° rotation), and was attached to the seed phantom as shown in Fig. 9. The seed phantom is comprised of twelve 5-mm-thick slabs, each having at least a hundred holes with 5 mm spacing. Any implant configuration with accurately known seed positions (in the FTRAC frame) can be created. Unfortunately there was about 0.5°–1° rotational error in the assembly of the attachment, leading to an error of about 0.5 mm in the ground truth estimates. Thus we expected additional error even before doing the experiments. The phantom is otherwise highly accurate. The seed density was kept constant at about 1.56 seed/cc. The number of seeds (and hence prostate size) was increased from 40 to 100 in steps of 15. For a given constellation, six images within a 20° cone around the AP axis were taken using an *uncalibrated Philips Integris V3000* fluoroscope. The obtained images were dewarped using the pin-cushion test. Accurate ground truth for matching was computed by utilizing the known 3D seed locations.

Matching was achieved using MARSHAL, followed by performance analysis. MARSHAL was evaluated for cases when two, two and a half, three, four, or more images are used. Robustness is further evaluated by using distorted images. For each implant, six images from various poses were

taken, and a number of combinations are created by choosing a different set of images to reconstruct the seeds. The results are averaged when displayed in Tables II and III.

**Two images:** The results for MARSHAL with two images, evaluated on 75 combinations (15 for each data set), are summarized on the left side of Table II. Two-image matching gives poor results, matching only about 2/3 of the seeds. The matched seeds reconstruct with an average RE of 1.17 mm, while the mismatched seeds reconstruct with an average error of 27.2 mm, the maximum being 241.8 mm. Though these results are completely unacceptable, MARSHAL provides the best possible solution that can exist with two images. Due to the various calibration and numerical errors, alternate constellations have significantly lesser error than the correct one. This is evident from the table where we can see that the average PE for mismatched seeds is only 0.22 mm (0.15 mm RA), but the seeds are deviated by a large value of 27.2 mm in 3D. Thus with two images, MARSHAL produces the best possible (least cost) matching that can exist.

**Two and a half images:** To *understand* better how information from a third image facilitates seed matching, partial information from a third image was added. Instead of all three, it runs only one bipartite matching, while still using Eq. (3) to compute  $C_{ij}$ . Since the third image is not completely used and this exercise is only toward our understanding, we say that only two and a half images are used. The results from 300 combinations ( $5 \times {}^6C_2 \times 4$ —the choice for the third image has four options) are shown on the right side of Table II. This little information from the third image is sufficient to correctly match 98.3% of the seeds. The mean RE for matched seeds is 0.88 mm, while that of mismatched seeds is 7.68 mm, the maximum for all combinations being 68.2 mm.

Thus when a third image is used, the number of good triplets decreases from  $O(N^{5/3})$  to  $O(N^{4/3})$ , providing a huge improvement in performance. Though the matching rate is

TABLE II. MARSHAL performance on phantom data. Two-image seed matching is completely erroneous (left), with over a quarter of the seeds being mismatched with high reconstruction error, yet low projection error and reconstruction accuracy. Adding just little information from a third image boosts performance (right). Most seeds match, though the reconstruction error is high for the few mismatched seeds. A total of 75 and 300 combinations were used, respectively.

		Number of seeds									
		Two images					Two and a half images				
		40	55	70	85	100	40	55	70	85	100
Matching rate (%)	Match	76.5	75.8	68.3	61.9	53.1	98.3	99.8	98.3	98.0	97.2
Reconstruction error (mm)		0.96	0.92	1.21	1.17	1.59	0.83	0.73	0.88	0.89	1.09
Projection error (mm)		0.19	0.12	0.15	0.15	0.19	0.20	0.16	0.17	0.18	0.24
Reconstruction error (mm)	Mismatch	23.9	27.8	28.2	23.0	33.0	2.35	13.2	4.32	8.38	10.2
Projection error (mm)		0.22	0.19	0.18	0.22	0.27	0.22	0.29	0.16	0.26	0.26
Reconstruction accuracy (mm)		0.15	0.13	0.13	0.15	0.19	0.13	0.09	0.12	0.16	0.15

good, RE for mismatched seeds is still rather high (note that PE for mismatched seeds is low). The reason is that in spite of being very close, the projection of the solution of the tripartite problem is not exactly the same as the solution of the bipartite problem. The proximity in the solution is evident from the low PE of 0.24 mm for the few mismatched seeds ( $<3\%$ ). Thus by including some combinatorial consistency checks, we should be able to match all the seeds. This intuition is realized fully in MARSHAL.

**Three images:** Averaged results from a total of 100 combinations ( $5 \times {}^6C_3$ ) are shown on the left side of Table III. 98.5% of the seeds matched perfectly. The 3D RE for matched seeds has a mean of 0.63 mm and STD 0.24 mm. The % of mismatched seeds is a strong upper-bound due to

the overlap problem (Sec. II H). This is evident from the RE for mismatched seeds, which has a mean value of 0.91 mm, only a little higher than correctly matched seeds. The average worst case error was 1.32 mm, the error being lower for smaller implants. As mentioned earlier, the ground truth of the seed locations has a small rotation error. To counter this, we computed *relative RE*, which removes any constant translation and rotation offset between the ground truth and reconstructed seeds, and measures any variation in the *shape* of the reconstructed implant. It can be seen that relative RE is significantly lower, standing at a mean value of 0.32 mm. Thus there is a constant shift of about 0.3 mm (due to the error in ground truth) for the whole implant, and an additional error of 0.32 mm for each seed. The absolute worst

TABLE III. Performance on phantom data. Using three images gives excellent results (left), with most of the seeds being matched. Mismatched seeds reconstruct with a low error. Using four images gives similar results (right). A total of 100 and 300 combinations were used, respectively.

		Number of seeds									
		Three images					Four images				
		40	55	70	85	100	40	55	70	85	100
Matching rate (%)	Match	97.6	100	98.0	98.5	98.5	99.2	100	98.7	98.9	99.0
Reconstruction error (mean)		0.60	0.48	0.63	0.70	0.76	0.59	0.48	0.63	0.69	0.75
Reconstruction error (STD)		0.21	0.23	0.25	0.24	0.25	0.21	0.23	0.25	0.24	0.25
Reconstruction error (mean)	Mismatch	0.73	...	0.76	0.84	1.30	0.67	...	0.65	0.64	1.14
Reconstruction error (worst)		1.18	...	1.03	1.10	1.96	0.91	...	0.86	0.77	1.58
Reconstruction error (relative)	All	0.28	0.29	0.35	0.30	0.39	0.28	0.29	0.35	0.29	0.38

case error was 3.29 mm, which seems to have a likelihood of about 1 in 75, as indicated from our data.

**Four or more images:** To seek further improvement, 300 combinations ( $5 \times {}^6C_3 \times 3$ ) using four images were run, the results of which are summarized on the right side of Table III. It can be seen that, on average, 99.2% of the seeds match perfectly. RE for matched seeds still has a mean of 0.63 mm and STD 0.24 mm. While RE of mismatched seeds has decreased to 0.78 mm and the average worst case error to 1.03 mm, the relative RE is still 0.32 mm. Thus we can see that adding a fourth image does not significantly improve the mismatched error, predominantly because three-image matching leaves very little scope for it. It should be noted that due to this redundancy, a fourth image is a great source for validation during the procedure.

Similar to using a fourth image, using a fifth or sixth image has only a little improvement in performance. Average results from 300 combinations indicate that using five images perfectly matches 99.2% of the seeds with the average worst case error dropping to 0.9 mm. Six images images match 99.4% with the average worst case error being 0.85 mm. Thus three images appears to achieve the correct balance between number of images used and accuracy obtained, though it should be noted that using more extra images would decrease RE and might also aid in the hidden seed problem.

It should be noted that %-matching implies a normalization by the number of seeds. Using four images, the 100-seed implant seems to perform similar to the 40-seed implant, matching around 99% of the total seeds. On the contrary, the 100-seed implant mismatches about 1 seed/implant while the 40-seed implant mismatches only 0.3 seeds/implant. In a similar manner, though the 55-seed implant matches 100% of the seeds, it does not imply that all 55-seed implants would always correctly match. Depending on a host of uncontrollable factors like errors in tracking, segmentation, noise, etc., an erroneous variation of about 1 seed/implant should typically be expected while studying any pattern.

#### D. Distorted images

MARSHAL addresses the problem globally as a combinatorial optimization problem, as compared to locally by using only epipolar constraints. Thus it is well-suited to evaluate the correspondences even in the presence of nonlinear effects like image distortion. Well-separated image positions (with distortion) were shown to match 98.5% of the seeds even on a 100-seed implant.<sup>67</sup> To further validate the robustness, we tested it for 700 combinations of distorted C-arm images. These combinations were chosen, so as to not restrict the images in any way, apart from being inside a 25° cone around the AP axis. Thus the images can be very close to each other and not well separated. Using four such images perfectly matching 93.2% of the seeds on an average, performing superior to three images that matched only 86.3%. Five and six images perfectly matched 95.4% and 96.2% of the seeds, respectively. The performance can be further improved by keeping the seed projections closer to the center of

the image. It can be observed from Fig. 9 that the seed phantom is not in the center of the image, and thus suffers from a considerable amount of distortion.

Thus, if the prostate size is small (<50 cc), then MARSHAL with three distorted images can match over 97% of the seeds match perfectly. This means that it flips about 2–3 seeds on average, correctly matching the rest. Alternatively, if the prostate is a little larger, then using more images that are well-spaced can still match about 98% of the seeds. Average RE for the whole implant stood high at 2.65 mm, while relative RE was low at 0.54 mm. The relative RE for the mismatched seeds was about 1.4 mm, while the absolute RE was 3.9 mm. Thus most seeds matched correctly, though the reconstructed implant is shifted in space by about 2 mm due to image distortion. The mismatched seeds have an additional error of 1 mm after the shift. Similar patterns were also observed on synthetic data in Sec. III B. Though we have not encountered any similar analysis in the literature, we conjecture that since those algorithms rely on epipolar lines that are distorted nonlinearly, they should be affected much more severely than MARSHAL.

It should be noted that though theoretically the reconstructed implant should be skewed, practically it is of the same shape. This is evident from the marginal increase in relative RE by only 0.23 mm. Thus the implant reconstructed from distorted images can be practically shifted by about 2 mm to overlap with the correct location. The reason for the absence of skew is that even though distortion is always nonlinear, the area that the implant occupies in the image is small. This results in an approximately linear shift of the seed coordinates in the image. Furthermore, 3D reconstruction from multiple images also reduces this error. As a consequence, the reconstructed implant is practically shifted in space in proportion to the amount of distortion.

Thus image distortion can potentially be either corrected using a constant function, or completely avoided if the amount is low. It should be noted that for this to work the implant should not be large and some mechanism needs to correct for the consistent shift in the x-ray to TRUS registration step. Moreover, additional images can provide significant improvement and reliability.

#### E. Run-time

MARSHAL has a run-time complexity of  $O(N^3)$  for any number of images, which makes it extremely fast even on large size implants. The unoptimized MATLAB code (Pentium4 3.2 GHz, Windows 2000, 1 Gbyte RAM) on a 100-seed implant with a seed density of 1.8 seeds/cc ran in 1.6 s. This run-time is typical of most tested implants, indicating that an optimized C++ implementation could in practice run in near real-time. Thus run-time for MARSHAL seems to be sufficient for intraoperative dose, especially when compared to the time required for automatic segmentation, which typically requires a few seconds. Though it was never observed in any of the thousands of test data sets, we point out that there can exist (theoretically) many long loops, which can potentially slow down MARSHAL.

#### IV. SHORTCOMINGS AND FUTURE WORK

The underlying assumption of this paper has been that all the seeds are segmented and their image coordinates known. Though excellent segmentation algorithms have been proposed,<sup>15,35,36</sup> they are capable of only segmenting seeds that are fully visible. Thus the coordinates of any hidden seeds (due to an overlap) will not be available, resulting in varying number of segmented seeds in each image. This is an issue of utmost importance and clinical practicality. While some researchers have reported the number of hidden seeds to be between 7% and 45% of the total seeds,<sup>35</sup> others have noticed only about 2%–3%<sup>68</sup> in their clinical data. Nevertheless, no solution can be clinically viable until it explicitly tackles this problem.

In some of the suggested solutions, fast-CARS was extended<sup>35</sup> to tackle this, but the algorithm reconstructed a greater number of seeds than were actually inserted. Another variant of fast-CARS was proposed<sup>68</sup> by ordering the seeds using the epipolar constraints. Unfortunately, the algorithm “based on epipolar geometry requires co-planar imaging to perform reconstruction” and “cannot reconstruct undetected seeds if they exist in the same search restriction band.” Significant works as they are, the problem merits further research. An intensity-based method using tomosynthesis<sup>69</sup> has also been proposed. Unfortunately in their current implementation, they require a rather large number of images to achieve accurate reconstruction. Another method based on Hough trajectories<sup>70</sup> has been proposed to solve this problem, but again requires a large number of images. Nevertheless, these methods could lead to promising results in the future.

Though MARSHAL does not yet tackle hidden and spurious seeds, we believe that the proposed mathematical framework for seed matching is comprehensive. It can be extended to solve these problems in a more formal framework of combinatorial optimization in contrast to the previously mentioned heuristic methodologies. This formal approach is expected to provide a better control on the behavior of the algorithm, potentially offering a decrease in dependence on rules that might change behavior from one data set to another and an increase in robustness. Moreover, the complete problem can be tackled as a whole (global optimization) as compared to looking at one seed or a group of seeds with heuristic rules (local optimization). Though it might be argued that the full problem should be broken down into independent subset of seeds with each subset solved separately, this might not be feasible in the generic case. If the tripartite problem could be broken down into a set of smaller min-cost tripartite problems, then it would no longer be NP-hard. Thus any solution should try to address the problem as a whole.

The basics of our approach will be to redefine the problem into a formal framework. Since we have proved in this publication that the global optimum is not attainable by any polynomial algorithm, we will project the full problem into multiple lower dimensional problems involving all the seeds. Each of these subproblems would be casted into the network

flow framework and the global optima for each lower dimensional problems would be achieved in polynomial time. The solutions to the subproblems would later be integrated together to create the final solution. The predominant difference from the current framework will be that instead of forced unitary flows, variable flows will be handled. Thus all the seeds will be handled in unison and solved globally. In addition, theoretically assured bounds on the performance of the algorithm are currently being worked out. A combination of the above-mentioned promises to lead to an implementation that could potentially be used clinically.

#### V. CONCLUSION

In contrast to previously proposed methods, we have formalized the seed-matching problem. We have proposed an algorithm, abbreviated as MARSHAL. It typically runs with  $O(N^3)$  time complexity with any number of images. Using three images, it matched perfectly over 98.5% of the seeds at a density of 1.56 seeds/cc. MARSHAL in combination with the FTRAC fiducial can reconstruct the seeds with a mean error of 0.63 mm and a STD of 0.24 mm. The mismatched seeds had a mean error of 0.9 mm, while the most mismatched seed in each implant had an average error of 1.32 mm. The worst error across all data sets was 3.29 mm and has a low likelihood of 1 in 75 implant cases. Moreover, relative reconstruction error was 0.32 mm, which reflects the deviation in shape of the reconstructed implant when compared to the true implant.

The algorithm performs well for distorted images too, matching over 97% of the seeds for small implants. The simulation experiments indicate that MARSHAL is not sensitive to (a) image separation, (b) seed density, (c) number of seeds, and (d) C-arm calibration. It is also acceptably robust to (a) segmentation, (b) C-arm pose, and (c) distortion. It can reconstruct an implant when three or more images are used, with a robustness, precision, and speed that promises to be sufficient to support intraoperative dosimetry in prostate brachytherapy. Though we validate it only for brachytherapy, MARSHAL is sufficiently generic to be used for establishing correspondences across many synergistic applications.

#### ACKNOWLEDGMENTS

This work has been financially supported by NIH 1R43CA099374-01 and NSF EEC-9731478. We are grateful to Professor Christian Scheideler for helping us formulate the problem. We also thank Stephen Kubinak, John Sofranko, Jane Kwietkowski, and Scott Borzillary for helping with the experiments.

<sup>a</sup>Electronic mail: jain@cs.jhu.edu

<sup>b</sup>Electronic mail: gabor@cs.jhu.edu

<sup>1</sup>A. Jemal, “Cancer statistics,” *Ca-Cancer J. Clin.* **54**(1), 8–29 (2004).

<sup>2</sup>R. Peschel and J. Colberg, “Surgery, brachytherapy, and external-beam radiotherapy for early prostate cancer,” *Lancet Oncol.* **4**, 233–241 (2003).

<sup>3</sup>J. Blasko, T. Mate, J. Sylvester, P. Grimm, and W. Cavanagh, “Brachytherapy for carcinoma of the prostate: Techniques, patient selection, and clinical outcomes,” *Semin. Radiat. Oncol.* **12**, 81–94 (2002).

<sup>4</sup>G. Merrick, W. Butler, J. Lief, and A. Dorsey, “Is brachytherapy comparable with radical prostatectomy and external-beam radiation for clini-



- cally localized prostate cancer?," *Tech. Urol.* **7**, 12–19 (2001).
- <sup>5</sup>S. Nag, J. Ciezki, R. Cormack, S. Doggett, K. DeWynngaert, G. Edmundson, R. Stock, N. Stone, Y. Yu, and M. Zelefsky, "Intraoperative planning and evaluation of permanent prostate brachytherapy: Report of the American Brachytherapy Society," *Int. J. Radiat. Oncol., Biol., Phys.* **51**, 1422–1430 (2001).
- <sup>6</sup>J. Xue, E. Gressen, and T. Jefferson, "Feasibility of trus-based prostate post-implant dosimetry," *AAPM Annual Meeting*, July 2004, p. Poster.
- <sup>7</sup>V. A. Dumane, M. Zaider, G. N. Cohen, and W. M. F. Worman, "Combined ultrasound-fluoroscopy approach to the intraoperative detection of seeds in prostate brachytherapy," *ASTRO Annual Meeting, Poster*, Oct 3–7 2004.
- <sup>8</sup>F. Mitri, P. Trompette, and J. Chapelon, "Improving the use of vibroacoustography for brachytherapy metal seed imaging: A feasibility study," *IEEE Trans. Med. Imaging* **23**, 1–6 (2004).
- <sup>9</sup>D. Holmes, B. Davis, C. Bruce, and R. Robb, "3d visualization, analysis, and treatment of the prostate using trans-urethral ultrasound," *Comput. Med. Imaging Graph.* **27**, 339–349 (2003).
- <sup>10</sup>D. French, J. Morris, M. Keyes, and S. E. Salcudean, "Real-time dosimetry for prostate brachytherapy using trus and fluoroscopy," *MICCAI 2004*, pp. 983–991.
- <sup>11</sup>H. Han, K. Wallner, G. Merrick, W. Butler, S. Sutlief, and J. Sylvester, "Prostate brachytherapy seed identification on post-implant trus images," *Med. Phys.* **30**, 898–900 (2003).
- <sup>12</sup>P. Kumar, R. Good, B. Epstein, M. Hussain, and F. Bartone, "Fluoroscopy guided transperineal percutaneous permanent 125iodine implantation of prostate cancer," *Radiat. Med.* **3**, 161–167 (1985).
- <sup>13</sup>B. Prestidge, J. Prete, T. Buchholz, J. Friedland, R. Stock, P. Grimm, and W. Bice, "A survey of current clinical practice of permanent prostate brachytherapy in the United States," *Int. J. Radiat. Oncol., Biol., Phys.* **15**, 461–465 (1998).
- <sup>14</sup>P. Ravindran, C. Lewis, J. Alphonsi, P. Lindsay, and D. D'Souza, "A c-arm based intra-operative dosimetry system for trans-rectal ultrasound guided prostate implant with 125I seeds," in *Proceedings of the 14th International Conference on the Use of Computers in Radiation Therapy*, 2004.
- <sup>15</sup>D. Todor, M. Zaider, G. Cohen, M. Worman, and M. Zelefsky, "Intraoperative dynamic dosimetry for prostate implants," *Phys. Med. Biol.* **48**, 1153–1171 (2003).
- <sup>16</sup>D. Todor, G. Cohen, H. Amols, and M. Zaider, "Operator-free, film-based 3d seed reconstruction in brachytherapy," *Phys. Med. Biol.* **47**, 2031–2048 (2002).
- <sup>17</sup>L. Gong, P. Cho, B. Han, K. Wallner, S. Sutlief, S. Pathak, D. Haynor, and Y. Kim, "Ultrasonography and fluoroscopic fusion for prostate brachytherapy dosimetry," *Int. J. Radiat. Oncol., Biol., Phys.* **1**, 1322–1330 (2002).
- <sup>18</sup>G. Grado, T. Larson, C. Balch, M. Grado, J. Collins, J. Kriegshauser, G. Swanson, R. Navickis, and M. Wilkes, "Actuarial disease-free survival after prostate cancer brachytherapy using interactive techniques with biplane ultrasound and fluoroscopic guidance," *Int. J. Radiat. Oncol., Biol., Phys.* **1**, 289–299 (1998).
- <sup>19</sup>G. Merrick, W. Butler, A. Dorsey, and H. Walbert, "Prostatic conformal brachytherapy: 125i/103pd postoperative dosimetric analysis," *Radiat. Oncol. Invest.* **5**, 305–313 (1997).
- <sup>20</sup>D. Nori and J. Moni, "Current issues in techniques of prostate brachytherapy," *Semin Surg. Oncol.* **13**, 444–453 (1997).
- <sup>21</sup>S. Nag, D. Scaperroth, R. Badalament, S. Hall, and J. Burgers, "Transperineal palladium 103 prostate brachytherapy: Analysis of morbidity and seed migration," *Urology* **45**, 87–92 (1995).
- <sup>22</sup>K. Wallner, S. Chiu-Tsao, J. Roy, V. Arterbery, W. Whitmore, S. Jain, B. Minsky, P. Russo, and Z. Fuks, "An improved method for computerized tomography-planned transperineal 125iodine prostate implants," *J. Urol. (Baltimore)* **146**, 90–95 (1991).
- <sup>23</sup>R. Hofstetter, M. Slomczykowski, M. Sati, and L. Nolte, "Fluoroscopy as an imaging means for computer-assisted surgical navigation," *Comput. Aided Surg.* **4**, 65–76 (1999).
- <sup>24</sup>J. Yao, R. H. Taylor, R. P. Goldberg, R. Kumar, A. Bzostek, V. R. Van, P. Kazanzides, and A. Guezic, "A c-arm fluoroscopy-guided progressive cut refinement strategy using a surgical robot," *Comput. Aided Surg.* **5**, 373–390 (2000).
- <sup>25</sup>A. Jain, T. Mustufa, Y. Zhou, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking (frac) fiducial," *Med. Phys.* **32**, 3185–3198 (2005).
- <sup>26</sup>A. Fung, "C-arm imaging for brachytherapy source reconstruction: Geometrical accuracy," *Med. Phys.* **29**, 724–726 (2002).
- <sup>27</sup>S. Narayanan and P. Cho, "3d seed reconstruction from unknown imaging geometry," in *Proceedings of the 14th International Conference on the Use of Computers in Radiation Therapy*, 2004, pp. 762–765.
- <sup>28</sup>M. Murphy and D. Todor, "Demonstration of a forward iterative method to reconstruct brachytherapy seed configurations from x-ray projections," *Phys. Med. Biol.* **50**, 2715–2737 (2005).
- <sup>29</sup>A. Jain, R. Kyon, and G. Fichtinger, "C-arm calibration—Is it really necessary?," *Medical Image Computing and Computer Assisted Interventions (MICCAI)*, 2005.
- <sup>30</sup>K. Wallner, J. Roy, M. Zelefsky, Z. Fuks, and L. Harrison, "Fluoroscopic visualization of the prostatic urethra to guide transperineal prostate implantation," *Int. J. Radiat. Oncol., Biol., Phys.* **1**, 863–867 (1994).
- <sup>31</sup>P. Roberson, V. Narayana, D. McShan, R. Winfield, and P. McLaughlin, "Source placement error for permanent implant of the prostate," *Med. Phys.* **24**, 251–257 (1997).
- <sup>32</sup>M. Baird, R. Holt, and T. Selby, "Improvement of transperineal implant dosimetry by intraoperative cystoscopic confirmation of prostate anatomy," *J. Urol. (Baltimore)* **164**, 406–410 (2000).
- <sup>33</sup>L. Archambault, L. Beaulieu, and D. Tubic, "Automatic post-implant needle reconstruction algorithm to characterize and improve implant robustness analyses," *Med. Phys.* **30**, 2897–2903 (2003).
- <sup>34</sup>M. Zhang, M. Zaider, M. Worman, and G. Cohen, "On the question of 3d seed reconstruction in prostate brachytherapy: The determination of x-ray source and film locations," *Phys. Med. Biol.* **49**, N335–345 (2004).
- <sup>35</sup>Y. Su, B. Davis, M. Herman, and R. Ra, "Prostate brachytherapy seed localization by analysis of multiple projections: Identifying and addressing the seed overlap problem," *Med. Phys.* **31**, 1277–1287 (2004).
- <sup>36</sup>D. Tubic, A. Zaccarin, J. Pouliot, and L. Beaulieu, "Automated seed detection and three-dimensional reconstruction. II. Reconstruction of permanent prostate implants using simulated annealing," *Med. Phys.* **28**, 2272–2279 (2001).
- <sup>37</sup>S. Lam, P. Cho, and R. Marks, "Prostate brachytherapy seed segmentation using spoke transform," *Proc. SPIE* **4322**, 1490–1500 (2001).
- <sup>38</sup>P. Cho, "Computerized segmentation of clustered seeds in prostate brachytherapy," *International Conference on the Use of Computers in Radiation Therapy*, 2000, pp. 105–107.
- <sup>39</sup>H. Amols and I. Rosen, "A three-film technique for reconstruction of radioactive seed implants," *Med. Phys.* **8**, 210–214 (1981).
- <sup>40</sup>P. Biggs and D. Kelley, "Geometric reconstruction of seed implants using a three-film technique," *Med. Phys.* **10**, 701–704 (1983).
- <sup>41</sup>M. Rosenthal and R. Nath, "An automatic seed identification technique for interstitial implants using three isocentric radiographs," *Med. Phys.* **10**, 475–479 (1983).
- <sup>42</sup>M. Altschuler, P. Findlay, and R. Epperson, "Rapid accurate, three-dimensional location of multiple seeds in implant radiotherapy treatment planning," *Phys. Med. Biol.* **28**, 1305–1308 (1983).
- <sup>43</sup>D. Jackson, "An automatic method for localizing radioactive seeds in implant dosimetry," *Med. Phys.* **10**, 370–372 (1983).
- <sup>44</sup>S. Li, G. Chen, C. Pelizzari, C. Reft, J. Roeske, and Y. Lu, "A new source localization algorithm with no requirement of one-to-one source correspondence between biplane radiographs," *Med. Phys.* **23**, 921–927 (1996).
- <sup>45</sup>R. L. Siddon and L. M. Chin, "Two-film brachytherapy reconstruction algorithm," *Med. Phys.* **12**, 77–83 (1985).
- <sup>46</sup>M. Altschuler and A. Kassae, "Automated matching of corresponding seed images of three simulator radiographs to allow 3d triangulation of implanted seeds," *Phys. Med. Biol.* **42**, 293–302 (1997).
- <sup>47</sup>S. Narayanan, P. Cho, and R. Marks, "Fast cross-projection algorithm for reconstruction of seeds in prostate brachytherapy," *Med. Phys.* **29**, 1572–1579 (2002).
- <sup>48</sup>K. Thornton, "Prostate brachytherapy seed localization by fluoroscopy/ultrasound fusion: Algorithms and analysis," *AAPM Annual Meeting*, July 2004.
- <sup>49</sup>H. Liu, G. Cheng, Y. Yu, R. Brasacchio, D. Rubens, J. Strang, L. Liao, and E. Messing, "Automatic localization of implanted seeds from post-implant ct images," *Phys. Med. Biol.* **48**, 1191–1203 (2003).
- <sup>50</sup>D. Dubois, W. J. Bice, and B. Prestige, "Ct and mri derived source localization error in a custom prostate phantom using automated image coregistration," *Med. Phys.* **28**, 2280–2284 (2001).
- <sup>51</sup>K. Shafique and M. Shah, "A non-iterative greedy algorithm for multi-frame point correspondence," *ICCV*, 2003, p. 6.

- <sup>52</sup>C. J. Veenman, M. J. T. Reinders, and E. Backer, "Resolving motion correspondence for densely moving points," *IEEE Trans. Pattern Anal. Mach. Intell.* **23**, 54–72 (2001). [Online.] Available: [citeseer.ist.psu.edu/veenman01resolving.html](http://citeseer.ist.psu.edu/veenman01resolving.html)
- <sup>53</sup>C. Veenman, M. Reinders, and E. Backer, "Establishing motion correspondence using extended temporal scope," *Artificial Intelligence* **145**, 227–243 (2003).
- <sup>54</sup>V. Salari and I. Sethi, "Feature point correspondence in the presence of occlusion," *Pattern Anal. Machine Intell.* **12**, 87–91 (1990).
- <sup>55</sup>Y. Cheng, R. Collins, A. Hanson, and E. Riseman, "Triangulation without correspondences," *ARPA94*, 1994, pp. II:993–1000. [Online.] Available: [citeseer.ist.psu.edu/6193.html](http://citeseer.ist.psu.edu/6193.html)
- <sup>56</sup>Y. Cheng, V. Wu, R. Collins, A. Hanson, and E. Riseman, "Maximum-weight bipartite matching technique and its application in image feature matching," 1996. [Online.] Available: [citeseer.ist.psu.edu/cheng96maximumweight.html](http://citeseer.ist.psu.edu/cheng96maximumweight.html)
- <sup>57</sup>*Artificial Intelligence: A Modern Approach*, Prentice Hall Series in Artificial Intelligence, edited by S. Russell and P. Norvig, 2nd ed., Chap. 24.5, pp. 885–892.
- <sup>58</sup>T. H. Cormen, C. E. Leiserson, and R. Rivest, *Introduction to Algorithms* 2nd ed. (MIT Press, 1990).
- <sup>59</sup>J. Edmonds and R. M. Karp, "Theoretical improvements in algorithmic efficiency for network flow problems," *J. ACM* **19**, 248–264 (1972).
- <sup>60</sup>A. Goldberg, "Network optimization library," <http://www.avglab.com/andrew/soft.html>.
- <sup>61</sup>Algorithmic Solutions Software GmbH, "Library of efficient data types and algorithms (leda)," [www.algorithmic-solutions.info/leda\\_guide/graph\\_algorithms/mincost\\_flow.html](http://www.algorithmic-solutions.info/leda_guide/graph_algorithms/mincost_flow.html).
- <sup>62</sup>R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, 1st ed. (Prentice Hall, Englewood Cliffs, NJ, 1993).
- <sup>63</sup>H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Res. Logistics Quart.* **2**, 83–97 (1955).
- <sup>64</sup>C. H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity* (Prentice-Hall, Englewood Cliffs, NJ, 1982).
- <sup>65</sup>G. Herman and A. Kuba, *Discrete Tomography: Foundations, Algorithms, and Applications* (Birkhauser, Boston, 1999).
- <sup>66</sup>K. Wallner, J. Blasko, and M. Dattoli, *Prostate Brachytherapy Made Complicated*, 2nd ed. (Smart Medicine, 2001).
- <sup>67</sup>A. Jain, Y. Zhou, T. Mustafa, E. C. Burdette, G. Chirikjian, and G. Fichtinger, "Matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL)," *SPIE Medical Imaging; Visualization, Image-Guided Procedures, and Display*, 2005.
- <sup>68</sup>S. Narayanan, P. Cho, and R. Marks, "Three-dimensional seed reconstruction from an incomplete data set for prostate brachytherapy," *Phys. Med. Biol.* **49**, 3483–3494 (2004).
- <sup>69</sup>I. Tutar, R. Managuli, V. Shamdasani, P. Cho, S. Pathak, and Y. Kim, "Tomosynthesis-based localization of radioactive seeds in prostate brachytherapy," *Med. Phys.* **30**, 3135–3142 (2003).
- <sup>70</sup>S. Lam, P. Cho, R. Marks, and S. Narayanan, "Three-dimensional seed reconstruction for prostate brachytherapy using hough trajectories," *Phys. Med. Biol.* **49**, 557–569 (2004).