Matching and Reconstruction of Brachytherapy Seeds using the Hungarian Algorithm (MARSHAL)

Ameet Kumar Jain[a], Yu Zhou[b], Tabish Mustufa[b], E. Clif Burdette[d], Gregory S. Chirikjian[b], and Gabor Fichtinger[a,c]

[a]Department of Computer Science, Johns Hopkins University
[b]Department of Mechanical Engineering, Johns Hopkins University
[c]Department of Radiology, Johns Hopkins University
[d]Computerized Medical Systems, Inc.

ABSTRACT

Purpose: Intraoperative dosimetric quality assurance in prostate brachytherapy critically depends on discerning the 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. **Method:** 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 network flow problem, which has salient features: (a) extensively studied exact solutions, (b) performance claims on the space-time complexity, (c) optimality bounds on the final solution. A fast implementation is realized using the Hungarian algorithm. **Results:** We prove that two images can correctly match only about 67% of the seeds, and that a third image renders the matching problem to be of non-polynomial complexity. We utilize the special structure of the problem and propose a pseudo-polynomial time algorithm. Using three images, MARSHAL achieved 100% 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 0.91 mm for incorrectly matched seeds. **Conclusion:** Both on synthetic data and in phantom experiments, matching rate and reconstruction accuracy were found to be sufficient for prostate brachytherapy. The algorithm is extendable to deal with arbitrary number of images without loss in speed or accuracy. The algorithm is sufficiently generic to be used for establishing correspondences across any choice of features in different imaging modalities.

Keywords: C-arm, Fluoroscopy, Seed Matching, Reconstruction, Prostate Brachytherapy, Radiation Planning

1. 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. For several decades, the definitive treatment for low risk prostate cancer was radical prostatectomy or external beam radiation therapy, but low dose rate permanent seed brachytherapy (shortly brachytherapy thereafter in this document) today can achieve virtually equivalent outcomes. The success of brachytherapy chiefly depends on our ability to tailor the therapeutic dose to the patient's individual anatomy. Prostate brachytherapy is almost exclusively performed with transrectal ultrasound (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 various methods, but even when meticulously hand-segmented, up to 25% of the seeds may remain hidden in ultrasound, which necessitates some other imaging method in intraoperative seed localization.

Several groups have published protocols and clinical outcomes favorably supporting intraoperative C-arm fluoroscopy. Mobile C-arms are ubiquitous in contemporary prostate brachytherapy, with approximately 60%of the practitioners using it for qualitative implant analysis in the operating room. The ability to reconstruct and register the implanted seeds (that are visible in fluoroscope) to soft tissue anatomy (that is visible in TRUS)

Send correspondence to E-mail: {jain, gabor}@cs.jhu.edu

Medical Imaging 2005: Visualization, Image-Guided Procedures, and Display, edited by Robert L. Galloway, Jr., Kevin R. Cleary, Proc. of SPIE Vol. 5744 (SPIE, Bellingham, WA, 2005) · 1605-7422/05/$15 · doi: 10.1117/12.595825

intraoperatively, would allow us to make immediate provisions for dosimetric deviations from the optimal implant plan. The three major obstacles we face towards intraoperative dosimetry are: (a) discerning the 3D pose of the fluoro images, (b) registering fluoro space to TRUS space, and (c) reconstruct the position of seeds from multiple fluoro images. Tracking un-encoded C-arms and registering fluoroscopic images to TRUS did not have a clinically practical solution, which we have solved by using the FTRAC fiducial.[1] Thus seed matching and reconstruction remains the missing link to intraoperative dosimetry, which is addressed here.

The matching problem is also prevalent in the computer vision community, where 2D points are tracked and reconstructed to compute motion.[2] Three-dimensional coordinates of the implanted seeds can be calculated from multiple X-ray images upon resolving the correspondence of seeds. Formalization of the seed-matching problem results in a high complexity search space of the order 10^{150} and 10^{300}, from 2 and 3 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 first step towards mathematical formalization came with the construction of a cost matrix,[3] where exhaustive matching gave the lowest cost solution, and required impractical computational resources. A greedy randomized algorithm,[4] tested with various cost metrics, was suggested to reduce the run-time. This method can give a different output for each run, and is not appropriate for clinical use. Fast-CARS[5] is another variant, which improved the computational complexity by doing the brute-force search on a subset instead of the whole space. It made the search significantly faster, yet it still ran in exponential time and global optimality was not assured. Fast-CARS was extended[6] to tackle hidden seeds, but reconstructed more number of seeds than were actually inserted. Another variant of fast-CARS was proposed[7] by ordering the seeds using the epipolar constraints, but it required all seeds to be identified on one image. Independently, a set of heuristic rules[8] were suggested for seed matching that attempted to reduce misclassifications. Simulated annealing[9] was proposed as an alternate technique to reach the global minimum. Another technique[10] 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. Moreover, none of these methods addressed the question of optimality of the solution.

2. MATERIALS AND METHODS

We convert seed-matching to network flow based combinatorial optimization. Our formulation has many salient features: (a) exact solutions studied extensively by the computer science community, (b) performance claims on the space-time complexity of the algorithm, (c) optimality bounds on the final solution, (d) guaranteed existence of a polynomial time solution for the global minima for seed-matching from 2 images, (e) proof of the non-existence of a polynomial time solution in case of more than 2 images, (f) derivation of a practical solution that can work near polynomial time on any number of images. We do not assume any prior information like inserted seed positions given by the dosimetric plan, the value of which is questionable due to intraoperative swelling and seed migration. In what follows, we assume that the seeds are 3D points and that their image locations are known, *i.e.* we do not solve automatic segmentation for which methods are already available.

2.1. Network flow based problem formalization

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 i^{th} seed in j^{th} image. We construct a directed network as shown in Figure 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 non-zero 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.

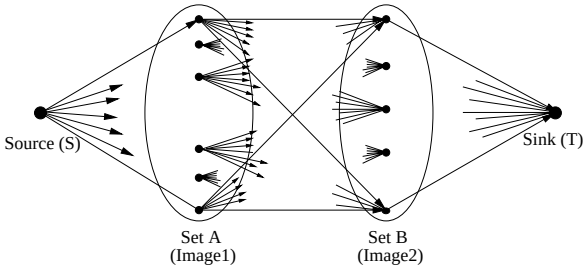

Figure 1. The seed matching problem can be converted to a network flow graph. The best possible matching reduces to finding out the maximum flow with minimum cost.

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 $\mathcal{C}_{ij}$. The cost $\mathcal{C}_{ij}$ represents the likelihood of seed s_{i1} matching seed s_{j2}, with the cost being 0 if they match perfectly and ∞ (infinity) if they don't match at all. Further details are provided in Section 2.4. Any feasible flow has a net cost associated with it, the value of which is $\sum_{i=1}^{n}\sum_{j=i}^{n}\mathcal{C}_{ij}f_{ij}$, where f_{ij} is the flow in link ij and $\mathcal{C}_{ij}$ 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

$$\begin{aligned} & min \sum_{i=1}^{N}\sum_{j=1}^{N}\mathcal{C}_{ij}f_{ij}, \\ & where \quad f_{ij} \in \{0,1\}; \\ & \sum_{j=1}^{N} f_{ij} = 1 \ \ \forall \ \ i \quad and \quad \sum_{i=1}^{N} f_{ij} = 1 \ \ \forall \ \ j \end{aligned} \tag{1}$$

The min-cost flow can be computed using cycle-cancelling, successive shortest path, primal-dual, out-of-kilter or relaxation algorithms. Though they are straightforward to use, they run in pseudo-polynomial time. The first weakly polynomial time algorithm was derived using an idea called scaling. 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. A comprehensive review of the above is available in the literature.[11] Hence the seed matching algorithm will have a polynomial run-time. Though strongly polynomial time algorithms are available, in practice the basic cycle-cancelling 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 theoretical run-times are the worst possible for generic networks.

2.2. 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. This is illustrated in Figure 2. 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 Equation (1). The assignment problem is solved in $O(N^3)$ run-time by using the Hungarian algorithm.[11, 12]

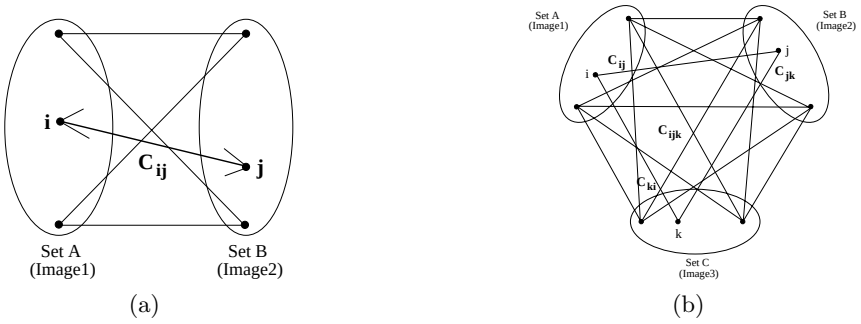

Figure 2. (a) Two image seed matching reduces to the bipartite problem. The best possible solution can be achieved in $O(N^3)$ time using the Hungarian algorithm. (b) The three-image seed matching problem reduces to the min-cost tripartite matching problem. No polynomial time solution is possible to compute the min-cost matching.

2.3. Seed Matching from 3 images

Due to projective imaging singularities, reconstructing brachytherapy seeds using discrete number of images is inherently inaccurate. Figure 2.3 (a) illustrates this intuition, where $Image_1$ and $Image_2$ cannot distinguish between the seed constellations A and B. A third image is required to remove this singularity. If the third image is coplanar, singular constellations still spring up, as illustrated in Figure 2.3 (B). Due to the small size of the prostate, the X-rays entering it are nearly parallel. Singularities though reduced, are not completely removed with non-coplanar images, as illustrated in Figure 2.3 (c). Some singular constellations cannot be resolved by four images either, which is achieved by putting together two parallelepipes along the fourth imaging direction. This process can construct singularities with any number of images. The likelihood of any of the above happening increases for large implants, especially since the seeds are inserted in parallel straight lines. Errors are further amplified with inaccuracies in segmentation, C-arm calibration and tracking. *Theoretically* seven good images can resolve all singularities, though *practically* three to four non-coplanar images should 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, as illustrated in Figure 2. Tripartite matching is similar to bipartite matching, except that it matches three sets instead of two. Multiple-image based seed matching reduces similarly to the multipartite matching problem. The bipartite problem is solvable in $O(N^3)$, while the tripartite and multipartite problems are NP-complete. Thus it is unlikely that there exists any polynomial time algorithm which can solve the tripartite problem. Hence we design our own algorithm that incorporates the physics of projective imaging, and can provide superior solutions.

The tripartite matching problem is formalised in Equation (2), where $\mathcal{C}_{ijk}$ is an $N \times N \times N$ cost-matrix. $\mathcal{C}_{ijk}$ in our case has a special structure. While a point in $Image_1$ constrains the projection in $Image_2$ to the epipolar line, both images constrain the projection in $Image_3$ to exactly one point. As a result $\mathcal{C}_{ij}$ being small did not imply that $< i, j >$ has to be a match. In contrast, a small $\mathcal{C}_{ijk}$ implies that $< i, j, k >$ is a good match. Statistically, two-image matching has an expected $O(N\sqrt{N})$ feasible doublets, while three image matching has just $O(N)$ feasible triplets. As the number of images increases, the number of feasible tuplets will rapidly converge to N, implying a decreasing number of feasible implant constellations. It should be observed that even having

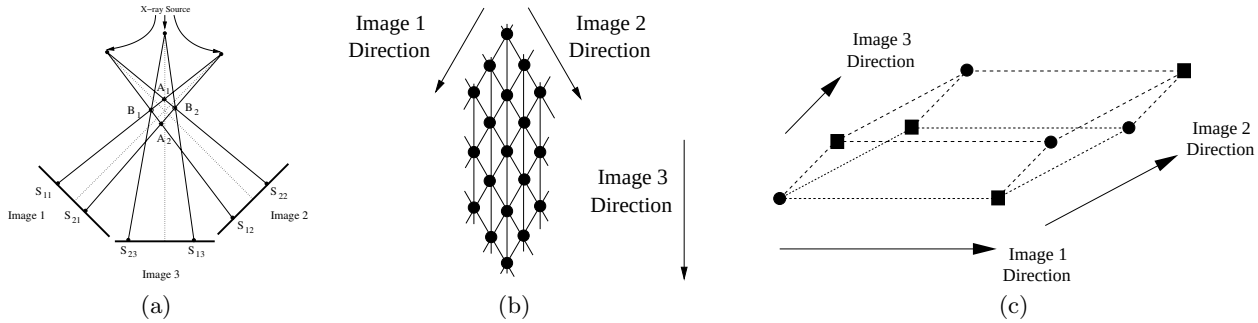

Figure 3. (a) Two images are inherently incapable of matching all the seeds. A third image is needed to resolve between the singular constellations A and B (b) Seed constellation singularity arising when three coplanar images are used. Multiple seed constellations will produce the same X-ray images. (c) Seed constellation singularity arising when three non-coplanar 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 even more images are used.

only $\alpha \times N$ tuplets, results in an exponential search space of $O(2^N \times (\alpha-1)^N)$ for large α and $O([\frac{2}{\alpha-1}]^{[\alpha-1]N})$ for small α. A formal proof for the non-existence of a polynomial-time algorithm can be shown by constructing a graph with all feasible tuplets as vertices, and compatibility between triplets as edges. The solution will be equivalent to finding a min-weight *clique* of size N, where a clique is a subgraph with edges between all the vertices. Finding a feasible clique itself is non-polynomial, let alone finding the clique with minimum cost.[13] Thus even with just $O(N)$ good triplets, an exact polynomial time algorithm is out of reach. A practical solution on the other hand, seems very possible.

$$
\begin{aligned}
& min \sum_{i=1}^{N}\sum_{j=1}^{N}\sum_{k=1}^{N} \mathcal{C}_{ijk} f_{ijk}, \\
& where \quad f_{ijk} \in \{0,1\}; \qquad \sum_{i=1}^{N}\sum_{j=1}^{N} f_{ijk} = 1 \ \forall \ k \ ; \\
& \sum_{j=1}^{N}\sum_{k=1}^{N} f_{ijk} = 1 \ \forall \ i \quad and \quad \sum_{k=1}^{N}\sum_{i=1}^{N} f_{ijk} = 1 \ \forall \ j
\end{aligned}
\tag{2}
$$

We propose a method for matching and reconstruction of seeds using the Hungarian algorithm (MARSHAL). Using constraints of projective imaging, it projects the tripartite problem into the closest possible bipartite problems, where the theoretical framework can guarantee optimality. The solutions of the bipartite matchings are *back-projected* t obtain a solution to the original problem. Hence it is a bridge between the theoretical framework and the inherent structure in the problem. The intuition behind MARSHAL can be summarized as: (a) tripartite matching can be projected into bipartite matchings; (b) s_{i1} matches s_{j2} only if they have a counter-part in the third image; (c) a *low* $\mathcal{C}_{ijk}$ forces $\mathcal{C}_{ij}$, $\mathcal{C}_{jk}$ *and* $\mathcal{C}_{ki}$ to be *low*. Thus if three *independent* bipartite matchings claim that s_{i1} matches s_{j2}, s_{j2} matches s_{k3} and s_{k3} matches s_{i1}, the independence between the matches insinuates $< i, j, k >$ as a correct triplet. The *projection* of the tripartite matching is achieved using appropriate cost matrices, explained in Section 2.4.

The flowchart for MARSHAL is shown in Figure 4. Using the respective cost matrices, for the three image pairs it obtains three independent bipartite matchings $M_1 < i, j, - >$, $M_2 < -, j, k >$ and $M_3 < i, -, k >$. Each bipartite matching has an $O(N^3)$ run-time and assures optimality. Loops are created next, *i.e.* if $< i_1, j_1, - >, < -, j_1, k_1 >, < i_2, -, k_1 >, < i_2, j_2, - > \ \ldots \ < i_1, -, k_m >$ are matchings, then $< i_1, j_1, k_1, i_2, j_2, k_2, \ldots, k_m, i_1 >$

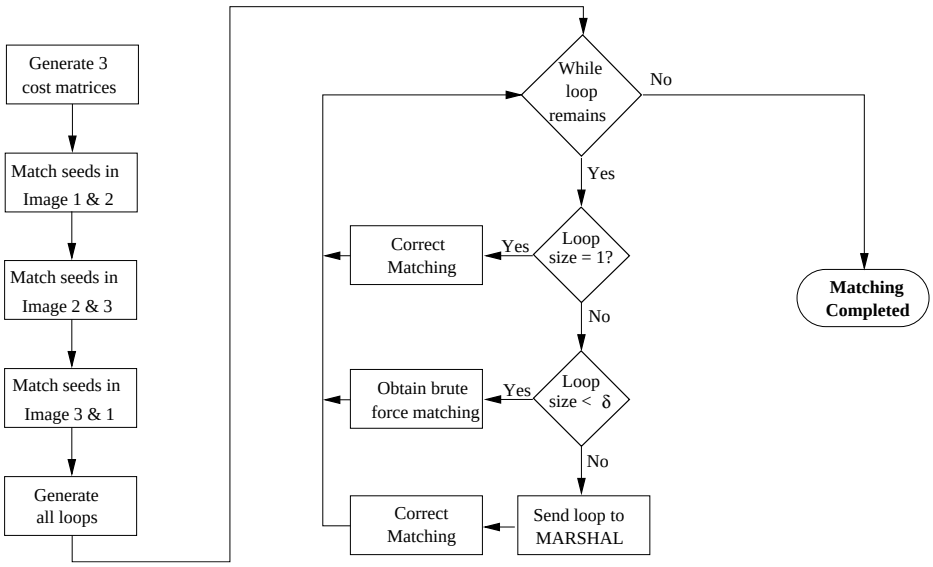

Figure 4. A flowchart explaining the working of MARSHAL. It runs in $O(N^3)$ time for *good datasets*.

is a loop of size m, each loop ending on the seed it started from. Now, each loop is a self-contained subset of seeds with no bearing on other loops. By a good projective cost matric, majority of the loops will have size 1. These can be declared matched.

Matchings sometimes get flipped, resulting in loops of size greater than 1. Since they have no bearing to the other loops, they only need to be matched among themselves to obtain the correct final answer. These are typically small loops, with the size never exceeding six. The optimal match can 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. Thus brute force search is done only for loops of size $m < \delta$, which is a pre-determined threshold based on the largest m that achieves a very fast (< 1 sec) run-time. Loops of size greater than δ 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, till everything is matched. Thus MARSHAL achieves the correct seed matching and still practically runs in $O(N^3)$.

2.4. Cost-metric

The two basic components of $\mathcal{C}_{ijk}$ are *reconstruction accuracy* (RA) and *projection error* (PE). $\mathcal{C}_{ijk}$ should represent the *consistency* among seeds s_{i1}, s_{j2} and s_{k3}, with the value being 0 for a perfect match. RA is the metric used by most researchers, measuring the accuracy with which the three projections can be reconstructed as a single 3D point. First, the equation of the three lines is computed that joins each projection to its respective X-ray source. Due to various errors these lines never intersect, creating a need for a symbolic intersection. A closed form solution for RA that minimizes the L_2 norm of the distance vector is used (Section 2.5). 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.

In our experience, PE fares significantly superior as a cost metric than RA. The reason is that PE directly measures the deviation from observed data, while RA is our own abstraction of consistency. Moreover, PE for each seed magnifies its RA, with different seeds having different magnifications. Thus PE is more sensitive to variations than RA, leading to better convergence. For example two seeds may have similar RA, but may have different PE depending on their placements. Thus any cost-metric that directly measures the deviation from the observation performs superior to a metric that does not. We have minimized the L_2 norm to compute the 3D

intersection, and used its PE as the metric. A combination of the above two techniques, gaining popularity in computer vision, is to choose the intersection so as to minimize the L_∞ norm of PE. Thus the most sensitive norm (L_∞) is minimized on the observations (seed locations on image), to compute the intersection. This metric would perform the best, though we have not used it due to a lack of a closed form solution.

$\mathcal{C}_{ij}$ needs to incorporate information from the third image to remove inherent two-image singularities. Since the bipartite matching should be the closest projection onto the tripartite matching,

$$\mathcal{C}_{ij} \quad = \quad min \quad \{\mathcal{C}_{ijk}/k = 1 \ldots N\} \tag{3}$$

Equation (3) says that s_{i1} and s_{j2} can be consistent if and only if there exists some s_{k3} such that $\mathcal{C}_{ijk}$ is small. Thus if i and j are matched, then they have a counterpart in the third image. This is similar to projecting the minimum value along the three axes. The computation would take $O(N^3)$ time. We modify it to run in $O(N^2)$ time. 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 PE calculated, which is assigned to $\mathcal{C}_{ij}$. Though this $\mathcal{C}_{ij}$ does not strongly satisfy equation (3), it weakly satisfies it. Any consistent $< i, j >$ will choose the correct k, while an inconsistent $< i, j >$ does not have any correct counter-part k, and will anyway be eliminated during the bipartite matching. Thus our $\mathcal{C}_{ij}$ exactly satisfies Equation (3) in all the relevant cases, and is close otherwise.

2.5. 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.

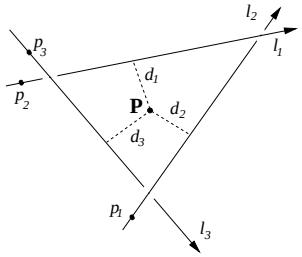

Figure 5. Three matching points, in general, do not intersect due to various errors in segmentation, C-arm pose and calibration. A point representing the symbolic intersection is needed.

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 Figure 5. 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 $(d_1, d_2, \ldots, d_m)$. In other words, we need to find a P such that it minimizes a function $\mathcal{F}$. Now, it can be easily computed that

$$\begin{aligned} \mathcal{F} = m \times RA^2 &= \sum_{i=1}^{m} d_i^2 \\ &= \sum_{i=1}^{m} ||(P - p_i) \times l_i||^2 \end{aligned} \tag{4}$$

$$\begin{aligned}\mathcal{F} &= \sum_{i=1}^{m}(P-p_i)^T \begin{bmatrix} b_i^2+c_i^2 & -a_ib_i & -a_ic_i \\ -a_ib_i & a_i^2+c_i^2 & -b_ic_i \\ -a_ic_i & -b_ic_i & a_i^2+b_i^2 \end{bmatrix} (P-p_i) \\ &= \sum_{i=1}^{m}(P-p_i)^T A_i (P-p_i)\end{aligned} \tag{5}$$

P minimizes $\mathcal{F}$. Hence we can write, $\frac{\partial \mathcal{F}}{\partial x} = \frac{\partial \mathcal{F}}{\partial y} = \frac{\partial \mathcal{F}}{\partial z} = 0$, which leads us to,

$$P = [\sum_{i=1}^{m} A_i]^{-1} \times [\sum_{i=1}^{m} A_i p_i] \tag{6}$$

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

2.6. Seed Matching from four images or more

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 very accurate, the run-time complexity is dependent on the number of images, which can make it slow, especially if extended to five or more images. Since the number of feasible m-tuplets rapidly decreases to N, we extend MARSHAL. The algorithm chooses 3 images at random as primary and the rest $(m-3)$ as secondary. Using the primary images, it runs exactly as described in Figure 4, except that each bipartite problem is a projection of the m-partite problem. In other words $\mathcal{C}_{ij}$ incorporates information from all the m images, as shown in Equation (7)

$$\mathcal{C}_{ij} = min\ \{\mathcal{C}_{i_1 i_2 \ldots i_m} / i_1 = i;\ i_2 = j\ and\ i_3, \ldots, i_m \in \{1, 2, \ldots, N\}\} \tag{7}$$

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

3. RESULTS AND DISCUSSION

Error analysis for seed matching is a delicate matter, typically using projection error (PE) and reconstruction accuracy (RA). In some cases, PE/RA are used to demarcate matched seeds from mismatched ones. As described in Section 2.3, this is invalid because incorrect seed constellations can produce exactly the same images. Thus any performance evaluation that is based on PE/RA is unreliable. The percentage of correctly matched seeds is the only accurate performance claim that can be made,and should be established by using a very accurate ground truth. An appropriate tool for measuring seed location accuracy is the 3D reconstruction error (RE), which is the 3D distance between the actual location of the seed and its reconstructed location. Moreover, it should be noted that it is more important to do the worst case analysis than the best case analysis. In this light, MARSHAL is evaluated using %-matching and RE.

3.1. Simulations

Validation was first performed on synthetic C-arm images of clinically realistic brachytherapy implant scenarios, where the reconstructed and actual seed locations were compared. Fifteen implants with the number of seeds varying from 60 to 150 in a 50 *cc* prostate were tested. The C-arm images were all contained inside a 15^o cone around the AP-axis. Table 1 shows the performance on this synthetic data. It can be noticed that two images are clearly insufficient for matching, correctly matching only 85.2% of the seeds. On the other hand, MARSHAL perfectly reconstructed 100% of the seeds when 3 images were used. The inherent imprecision in the two-image case is indicated by the low PE and high RE for the unmatched seeds. Also, a 120-seed implant was tested with 24 permutations of 3 images, each with 15^o separation. All seeds were matched correctly. Both the above experiments indicate that MARSHAL is not sensitive to the number of seeds and to the choice of C-arm poses.

Table 1. 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									
		2 Images						3 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	**Match**	0.18	0.15	0.18	0.20	0.21		0.07	0.07	0.07	0.07	0.07
Reconstruction error	**Mismatch**	22.1	21.4	22.1	23.6	21.9		-	-	-	-	-
Projection error	**Mismatch**	0.03	0.04	0.04	0.04	0.04		-	-	-	-	-

3.2. Phantom experiments

After the simulations, experiments were conducted on a precisely fabricated seed phantom, constructed using Delrin. The FTRAC fiducial[1] was used to track the C-arm (accuracy of 0.56 *mm* translation and 0.33^o rotation), and was attached to the seed phantom as shown in Figure 6. The seed phantom comprises of twelve 5 *mm* thick slabs, each having atleast 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 $0.5^o - 1^o$ 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*. For a given constellation, 6 images within a 20^o cone around the AP-axis were taken using a *Philips BV 3000* fluoroscope. Though error-prone, the manufacturer supplied fluoroscope parameters were used for reconstruction. The obtained images were dewarped using the pin-cushion test. By utilising the known 3D seed locations, seeds in all the images were hand-segmented to establish correspondence and the most accurate ground truth possible. Matching was achieved using MARSHAL, followed by performance analysis. MARSHAL was evaluated various number of images. Robustness is further evaluated by using distorted images. Table 2 sums up the result.

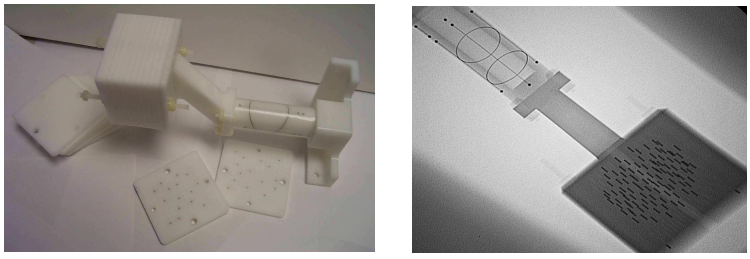

(a) (b)

Figure 6. (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

2 Images: The results for MARSHAL with two images, evaluated on 75 combinations, are summarized on the left side of Table 2. Two-image matching gives poor results, matching only about 2/3 of the seeds. The

Table 2. MARSHAL performance on phantom data. 2-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. Using three images gives excellent results (right), with most of the seeds being matched. Mismatched seeds reconstruct with a low error. A total of 75 and 100 combinations were used respectively.

		Number of Seeds									
		2 Images					3 Images				
		40	**55**	**70**	**85**	**100**	**40**	**55**	**70**	**85**	**100**
Matching Rate (%)	**Match**	76.5	75.8	68.3	61.9	53.1	97.6	100	98.0	98.5	98.5
Reconstruction Error (mm)		0.96	0.92	1.21	1.17	1.59	0.60	0.48	0.63	0.70	0.76
Reconstruction Error (mm)	**Mismatch**	23.9	27.8	28.2	23.0	33.0	0.73	-	0.76	0.84	1.30
Projection Error (mm)		0.22	0.19	0.18	0.22	0.27	0.31	-	0.27	0.45	0.46
Reconstruction Error (worst)		133	144	206	150	242	1.18	-	1.03	1.10	1.96
Reconstruction Error (relative)	**All**	6.72	8.00	10.0	9.91	16.8	0.28	0.29	0.35	0.30	0.39

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 only 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 implant constellation, which produces exactly the same C-arm images as the correct constellation.

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 a single bipartite matching, while still using Equation (3) to compute $\mathcal{C}_{ij}$. 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*. Thus when a third image is used, the number of good triplets decreases from $O(N\sqrt{N})$ to $O(N)$, providing a huge improvement in performance. By including some combinatorial consistency checks, we will be able to match all the seeds. This intuition is realised fully in MARSHAL.

3 Images: Averaged results from a total of 100 combinations are shown on the right side of Table 2. 98.5% of the seeds match perfectly. The 3D RE for matched seeds has a mean of 0.63 *mm* and STD 0.24 *mm*. Only 1.5% of the seeds are mismatched, which appears to be a very strong upper-bound due to the overlap problem. As illustrated in Figure 7, both matchings **A** and **B** give the correct 3D seed coordinates, though only matching **A** is theoretically correct. Owing to our rigid error analysis criterion, we declare **B** as an incorrect matching. 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* for the whole implant, and an additional error of 0.32 *mm* for each seed.

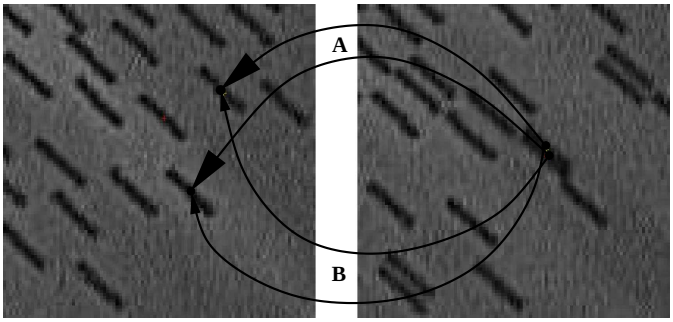

Figure 7. Two seeds nearly overlap in the right image. Though both matchings A and B reconstruct the seeds in the correct location, only one of them is theoretically correct. We define the other as mismatched, so as to make the seed matching error analysis extremely rigorous.

4 or more Images: To seek further improvement, 300 combinations using four images were run. 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. Similarly, using a fifth or sixth image has little improvement in performance. 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. When three images were used, the absolute worst case error was 3.29 mm, which has an occurrence likelihood of 1 in 300.

Run-time: The run-time for MARSHAL is very fast, unless there are lots of long loops to be resolved, in which case the brute-force computations within the long loops may slow it down by a few seconds. An un-optimized MATLAB code (Pentium4 2.4GHz, Windows 2000, 512MB RAM) typically runs in about 10 s, indicating that an optimized C++ implementation would take about a second to run. This should be sufficient for any intraoperative dosimetry based procedures.

3.3. Distorted Images

To further validate the robustness, we tested it for the 800 combinations without dewarping the C-arm images. Using four 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 match 95.4% and 96.2% perfectly. Average RE stood high at 2.5 mm, while relative RE is low at 0.5 mm. Thus most seeds matched correctly, though the reconstructed implant is shifted in space by 2 mm due to the warp. If the prostate size is small (50 cc), over 97% of the seeds match perfectly. This implies that distortion correction can be avoided when the implant size is small and if the consistent shift can be corrected during the X-ray to TRUS registration step. Though the above data is averaged over randomly chosen images, choosing three well separated good images can match 98.5% seeds even on a 100-seed implant.[14] Moreover, a fifth image can provide significant improvement and reliability.

4. CONCLUSION AND FUTURE WORK

In contrast to previously proposed methods, we have formalised the seed matching problem and proposed a robust algorithm, abbreviated as MARSHAL. It appears to be sufficiently robust for brachytherapy implant reconstruction, matching 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 worst from 300 combinations being 3.29 mm. Moreover, relative reconstruction error is 0.32 mm, and the algorithm performs well for warped images too. The simulation experiments indicate that it is not sensitive to the image separation or the number of seeds. Thus it can provide near perfect reconstruction of an implant when three of more images are used, with a robustness, precision, and speed that promises to be sufficient for intraoperative quality control in prostate brachytherapy.

Though, presently we have not tackled hidden and spuriously segmented seeds, we realize that it is an issue of utmost importance and practicality. The mathematical framework for seed matching that we have proposed is comprehensive, and hence can be extended to solve these problems using non-heuristic methodologies. This work is currently in progress. Furthermore, an algorithm to alleviate automated seed segmentation is being developed. A combination of all these promises to lead to an implementation that can be dependably used in brachytherapy. Though we validate it for brachytherapy, MARSHAL is sufficiently generic to be used for establishing correspondences across any choice of features, among images of the same or different imaging modalities.

5. ACKNOWLEDGMENTS

This work has been financially supported by NIH 1R43CA099374-01 and NSF EEC-9731478, and co-sponsored by Computerized Medical Systems. We are grateful for the thoughtful advice of Professor Russell H Taylor and Gregory D Hager. We also thank Stephen Kubinak, John Sofranko, Jane Kwietkowski and Scott Borzillary for helping with the experiments.

REFERENCES

1. A. K. Jain, T. Mustufa, Y. Zhou, E. C. Burdette, G. S. Chirikjian, and G. Fichtinger, "A robust fluoroscope tracking fiducial," in *SPIE medical imaging*, **accepted**, Feb 2005.
2. C. Veenman, M. Reinders, and E. Backer, "Establishing motion correspondence using extended temporal scope," *AI* **145**, pp. 227–243, April 2003.
3. R. L. Siddon and L. M. Chin, "Two-film brachytherapy reconstruction algorithm," *Med Phys* **12(1)**, pp. 77–83, Jan 1985.
4. M. Altschuler and A. Kassaee, "Automated matching of corresponding seed images of three simulator radiographs to allow 3d triangulation of implanted seeds," *Phys Med Biol* **42(2)**, pp. 293–302, Feb 1997.
5. S. Narayanan, P. Cho, and R. Marks, "Fast cross-projection algorithm for reconstruction of seeds in prostate brachytherapy," *Med Phys* **29(7)**, pp. 1572–9, Jul 2002.
6. 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(5)**, pp. 1277–87, May 2004.
7. S. Narayanan, P. Cho, and R. Marks, "Three-dimensional seed reconstruction from an incomplete data set for prostate brachytherapy," *Phys Med Biol* **49(15)**, pp. 3483–94, Aug 2004.
8. D. Todor, G. Cohen, H. Amols, and M. Zaider, "Operator-free, film-based 3d seed reconstruction in brachytherapy," *Phys Med Biol* **47(12)**, pp. 2031–48, Jun 2002.
9. D. Tubic, A. Zaccarin, B. L, and J. Pouliot, "Automated seed detection and three-dimensional reconstruction. ii. reconstruction of permanent prostate implants using simulated annealing," *Med Phys* **28(11)**, pp. 2272–9, Nov 2001.
10. M. Murphy and D. Todor, "Iterative reconstruction of brachytherapy seed configurations from x-ray projection images," in *AAPM annual meeting*, July 2004.
11. R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall; 1 edition, 1993.
12. H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logist. Quart* **2**, pp. 83–97, 1955.
13. T. H. Cormen, C. E. Leiserson, and R. Rivest, *Introduction to Algorithms*, MIT Press/McGraw-Hill, 1990.
14. A. K. Jain, Y. Zhou, T. Mustufa, E. C. Burdette, G. S. Chirikjian, and G. Fichtinger, "Brachytherapy seed reconstruction from fluoroscopic images using network flow algorithms," in *ASTRO annual meeting, Poster*, Oct 3-7 2004.