Issues in the Metrology of Geometric Tolerancing

Chee Yap and Ee-Chien Chang,
Courant Institute of Mathematical Sciences, New York University,
251, Mercer Street, New York, NY 10012, USA

The classification problem is fundamental in the metrology of Geometric Tolerancing: given a part B, is it within some tolerance Φ? The "standard methodology" for the classification problem comprises three steps: sampling σ, algorithmic step α and policy decision π. The issues we examine are posed around this methodology, using a simple 1-dimensional problem as case study.

August 14, 1996

1 Introduction

Geometric tolerancing is concerned with the specification of geometric shapes for use in the manufacture of mechanical parts. Since manufacturing processes are inherently imprecise, it is imperative that such geometric designs be accompanied by tolerance specifications. Figure 1 is a simple example of such a specification. Evidently, the ideal object indicated there is a 3×5

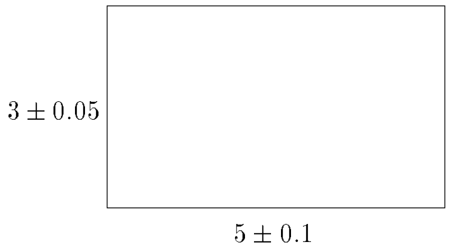

Figure 1: *Tolerancing a rectangle.*

rectangle, with some amount of deviation to be tolerated in its manufacture. There is a naive interpretation of this symbology, namely, "*any manufactured rectangle of dimension $L \times H$ where $|L-5| \leq 0.1$ and $|H-3| \leq 0.05$ is considered within tolerance*". This is the **parametric tolerancing** semantics. The problem with this approach is that no manufactured planar object[1] will be a rectangle or even a quadrilateral, but rather some planar shape with an uneven boundary. Following Requicha [9], we believe that a rigorous approach can be based on the **zone tolerancing semantics**. This interprets the symbology of figure 1 to specify some region called the **(tolerance) zone** and declare any body with boundary within this zone as "within tolerance".

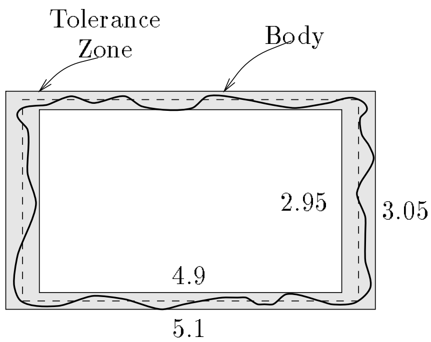

Figure 2: *Zone Semantics.*

Figure 2 illustrates one interpretation of the zone implicit in figure 1: it is a region bounded by two "aligned" rectangles with the indicated dimensions. We remark that the actual language for tolerance specifications used by professional engineers is highly intricate, involving both planar figures and textual symbols (e.g., [14, 8]) and requiring careful mathematical interpretation [13]. See [6] for an exposition. The other re-

[1] Parametric tolerancing may be useful in non-manufacturing applications. For instance, in a window-based graphical interface system, we might wish to tolerance the window sizes using this approach.

mark is that the semantics in the ANSI standard [14, 8] is basically a zone approach.

Having toleranced a geometric design, suppose we now hold in hand an actual (physical) object that is supposed to be an instance of the design. We need to check if this instance is within tolerance. The **metrology of geometric tolerancing** (or simply "tolerancing metrology") is the study of this measurement process[2]. Hopp [2] calls this subject "computational metrology". The use of **coordinate measuring systems** or **machines** (CMS or CMM) is considered the state-of-the-art in the practice of tolerancing metrology. Each CMS is basically a highly accurate mechanical probe, coupled to computational power. As such a CMS can be programmed to automatically take probes and perform computations.

We use the term "Geometric Tolerancing and Metrology" (GT&M) to refer to the combination of geometric tolerancing with computational metrology. The latter is by far the less understood half of GT&M. The present paper attempts to expose some basic metrological issues.

Importance of GT&M. Intellectually, GT&M is a cornerstone of modern high precision engineering. But this subject also has widespread implications in practice because achieving high precision in manufacturing is a major component of manufacturing costs. Studies have shown that manufacturing tolerances decrease linearly with time (on a semi-log scale), in every category of manufacturing, from large-scale manufacture such as shipbuilding to micron-level manufacture such as computer chips. Hence such issues are expected to become more acute with time.

A more comprehensive view of the "cost of a part" must look beyond the manufacturing cost. The "lifetime cost" of a manufactured part must include its reliability over its life-time. Reliability is positively correlated with tight tolerances, as is evident in the reliability of parts such as engine pistons. Hence we see a tradeoff between reliability versus manufacturing cost.

[2]Note that "metrology" is the science of measurement and as such is much wider than its applications to geometric tolerancing.

Just as manufacturing is inherently imprecise, metrology is also inherently imprecise. This has implications be making policies for accepting or rejecting manufactured parts: there is a cost for rejecting good parts as well as a cost for accepting bad parts. For instance, machining a high-precision part such as a gear for a jet engine may cost up to $50K. Without solid metrological principles, we may be hard pressed between making a $50K mistake (if we wrongly reject such a gear) versus making a graver mistake (if we wrongly accept such a gear).

Current tolerancing metrology practice often appears ad hoc. For instance, a subcontractor for a car manufacturer is to produce aluminum space-frames which are basically tubes that are bent to some specified shape. The contractual agreement calls for coordinate measuring machine (CMM) measurements at 64 locations in each space-frame, with the criteria is that each measurement q_i is within 1mm from given ideal measurement p_i: $\|q_i - p_i\| \leq 1$ $(i = 1, \ldots, 64)$. This turns out to be remarkably hard to achieve. It is also unclear if this contractual bases of performance is the most appropriate.

For overviews of current GT&M, we refer to [12, 10, 1]. A recent report [4] describes the economic impact of CMM research carried out at the National Institute of Standards and Technology (NIST) from 1975-85; it is also a good background for CMM research. See [15] for the connection to computational geometry and exact computation.

2 Standard Methodology in Classification

In metrology, our original datum is a **body** $B \in \mathbb{R}^n$ $(n = 1, 2, 3)$. For our purposes, define a body B to be any compact, simply connected subset of $\mathbb{R}^n$. The boundary of B is denoted ∂B. For each toleranced geometric design D, we obtain a corresponding set Φ_D comprising all bodies B that are "within tolerance". Figure 1(a) is an example of D. Viewing $\Phi_D(B)$ as a predicate, we have the corresponding **classification problem** for D: *given* B, *does* $\Phi_D(B)$ *hold?* The following three steps constitute the **standard methodology** for classification problems:

I. Measurement Step: Make some measurements S on the given body B. Let μ denote this measurement procedure,

$$\mu : B \mapsto S.$$

II. Algorithmic Step: From S, we compute a small set V of "characteristic values". Let α denote the algorithm for this,

$$\alpha : S \mapsto V.$$

III. Policy Step: Using V, we finally make a decision to accept or reject the body B. Let this decision be based on a policy π,

$$\pi : V \mapsto \{accept, reject\}.$$

Let us illustrate this methodology for the well-known "roundness problem" (e.g., [11]). The **(annular) width** of a body $B \subseteq \mathbb{R}^2$,

$$\texttt{width}(B)$$

is defined to be the minimum width $R-r$ of an annulus with radii (R, r) that contains the boundary of B. If S is a finite set, $\texttt{width}(S)$ is similarly defined. Suppose that we are interested in the classification predicate

$$\Phi(B) : \text{``}\texttt{width}(B) \leq 0.1''.$$

The standard methodology (μ, α, π) for this classification problem is typically as follows:

- μ: Probe B to get n sample points $S \subseteq \partial B$.
- α: Compute $\texttt{width}(S)$.
- π: Accept iff $\texttt{width}(S) \leq 0.1$.

The following questions arise:

(i) Accuracy: How shall n be chosen? Intuitively, it is clear that as n grows, the accuracy improves.

(ii) Sampling strategy: How shall S be obtained? In practice, the sampling strategies are ad hoc and unanalyzed.

(iii) Validity of substitution: The above policy apparently treats $\texttt{width}(S)$ as a substitute for $\texttt{width}(B)$. Why is this appropriate — what is their relation?

(iv) Policy analysis: Wrong decisions are inevitable in any policy (recall that metrology is inherently imprecise). How should we hedge our bets concerning making **erroneous rejections** versus **erroneous acceptances**? Borrowing a terminology of hypothesis testing in statistics, we may refer to these two kinds of errors as **type I** and **type II** errors, respectively.

(v) Integration: The connections between μ, α, π are often not exploited. In [7], an integrated approach to (μ, α, π) is given for the problem of classifying roundness of convex polygons.

We remark that often computational geometers simply address the algorithmic part (α) of this methodology. But from the metrological perspective, this is not the most significant aspect. We believe that computational geometry can contribute to the entire process.

3 A Simple Probe Model

To expose some issues, consider a toy problem in one dimension. Let $B \subseteq \mathbb{R}^1$ be a closed line segment of some unknown length L. The classification problem we address is this: does L satisfy the bounds

$$1 - \delta \leq L \leq 1 + \delta \tag{1}$$

where $0 < \delta < 1$ is fixed. We also write $L \in [1 \pm \delta]$ instead of (1). The probe model we consider is the **Δ-grid** for some $\Delta > 0$. It is comprised of infinitely many point "sensors" that are uniformly spaced Δ-apart along the entire real line. Figure 3 illustrates the Δ-grid. To make a probe, we simply place B on the grid and we obtain as **probe result** the number n_0 of sensors that are covered by B. Note that we have no control over how to place B, so that n_0 is not a deterministic function of L. In figure 3, the value of n_0 is 5.

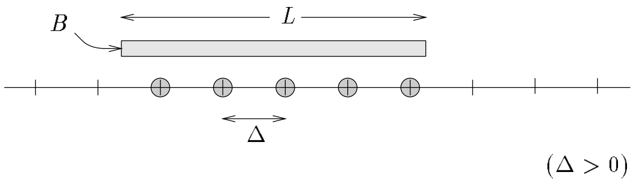

Figure 3: Δ-*grid.*

Nevertheless, we shall only consider[3] **deterministic** policies based on n_0.

E.g., A simple policy might be as follows: accept B iff $n_0 = \lfloor 1/\Delta \rceil$. Here $\lfloor 1/\Delta \rceil$ denotes the rounding of $1/\Delta$ to the closest integer. How good is this policy?

We classify policies into three basic types: (1) A policy is **conservative** if it never accepts a bad L. (2) It is **liberal** if it never rejects a good L. A conservative policy is **trivial** if it rejects every input; a liberal policy is likewise trivial if it accepts every input. The basic deduction from our probe result n_0 is the following:

$$(n_0 - 1)\Delta \le L < (n_0 + 1)\Delta.$$

It follows that a non-trivial conservative policy exists iff there exist $n_*, n^* \in \mathbb{N}$, $1 \le n_* \le n^*$, such that

$$[(n_* - 1)\Delta, (n^* + 1)\Delta] \subseteq [1 \pm \delta].$$

Choosing $n^* - n_*$ as large as possible, we have the "best" conservative policy: *accept iff* $n_0 \in [n_*, n^*]$. As corollary:

(i) If $\Delta > \delta$, no non-trivial conservative policy exists.

(ii) If $\Delta \le \delta/2$, non-trivial conservative policies exist.

(iii) If $\Delta \le \delta$ and Δ is an Egyptian fraction, non-trivial conservative policies exist.

Recall that an Egyptian fraction is one whose reciprocal is an integer. On the other hand, note that non-trivial liberal policies always exist. Indeed, (i)-(iii) are subsumed by the following:

[3] But our policies are necessarily nondeterministic when viewed as functions of L.

Lemma 1 *Let* $\varepsilon_0 = \lceil 1/\Delta \rceil - (1/\Delta)$, $\varepsilon_1 = (1/\Delta) - \lfloor 1/\Delta \rfloor$, *and* $\varepsilon = \min\{\varepsilon_0, \varepsilon_1\}$. *Then a non-trivial conservative policy exists iff* $\delta \ge (1+\varepsilon)\Delta$.

Except for the "perfect policy" (which does not exist under our probe model), note that conservative and liberal policies are mutually incompatible. Conservative (resp., liberal) policies may be useful when the cost of wrong acceptance (resp., rejection) is extremely high. But usually, a tradeoff is more useful. We therefore formulate another class of policies: let $0 \le \delta_* \le \delta^* \le \infty$. A policy is δ_***-neoliberal** if $|L - 1| < \delta_*$ implies that it surely accepts. A policy is δ^***-neoconservative** if $|L - 1| > \delta^*$ implies that it surely rejects. Say a policy is a (δ_*, δ^*)**-policy** if it is δ_*-neoliberal and δ^*-neoconservative. We say the (δ_*, δ^*)-policy is **compatible** with the classification of equation (1) if

$$\delta_* \le \delta \le \delta^*.$$

Note that a compatible (δ_*, δ^*)-policy is not constrained to behave in any particular way on inputs that satisfy $\delta_* \le |L - 1| \le \delta^*$: it may behave conservatively or liberally. Also, a non-trivial conservative policy may not be δ_*-neoliberal for any $\delta_* > 0$.

Lemma 2 (δ_*, δ^*)-
Policies exist iff (a) $\lceil (1-\delta^*)/\Delta \rceil < \lfloor (1-\delta_*)/\Delta \rfloor$, *and (b)* $\lceil (1+\delta_*)/\Delta \rceil < \lfloor (1+\delta^*)/\Delta \rfloor$. *Equivalently, such policies exist iff* $\delta^* - \delta_* \ge (1+\varepsilon)\Delta$ *where*

$$\varepsilon = \max\{\frac{1-\delta_*}{\Delta} - \left\lfloor \frac{1-\delta_*}{\Delta} \right\rfloor, \left\lceil \frac{1+\delta_*}{\Delta} \right\rceil - \frac{1+\delta_*}{\Delta}\}$$

EXPLICIT POLICIES. For $\alpha, \beta \ge 0$, the **policy** (α, β) is as follows: *Accept iff* $1 \in [(n_0 - \alpha)\Delta, (n_0 + \beta)\Delta]$. Not all (δ_*, δ^*)-policies can be achieved by explicit policies. The following shows some attainable bounds.

Lemma 3 *Fix* $\alpha, \beta \ge 0$. *For all* Δ, *the policy* (α, β) *is a* (δ_*, δ^*)*-policy where*

$$\begin{aligned} \delta_* &= \Delta \cdot \min_0\{\alpha - 1, \beta - 1\}, \\ \delta^* &= \Delta \cdot \max\{\alpha + 1, \beta + 1\}. \end{aligned} \qquad (2)$$

The notation $\min_0\{x, y\}$ *here refers to* $\max\{0, \min\{x, y\}\}$. *This is the best possible* *in the sense that there exists Δ_* (resp., Δ^*) for which this choice of δ_* (resp., δ^*) is optimal.*

Proof. We know that $L \in [(n_0 \pm 1)\Delta]$ and we accept iff $1 \in [(n_0 - \alpha)\Delta, (n_0 + \beta)\Delta]$. We want to determine a suitable δ_*, δ^*. Suppose $\alpha < 1$. We must show that $\delta_* = 0$ is "best possible" choice. For any $\varepsilon > 0$, there exists Δ such that the policy (α, β) rejects an input of length L where $|L - 1| = \varepsilon$: we simply choose Δ to satisfy $1 = (n_0 - \alpha)\Delta - \varepsilon$. Hence $\delta_* = 0$ is best possible in the sense of the lemma. Similarly if $\beta < 1$ then $\delta_* = 0$ is the best possible. On the other hand, if $\alpha \geq 1$ and $\beta \geq 1$, then $|L - 1| \leq \Delta \cdot \min\{\alpha - 1, \beta - 1\}$ implies that policy (α, β) will surely accept L. This proves the choice $\delta_* = \Delta \cdot \min\{\alpha - 1, \beta - 1\}$ is valid. It is the best possible in the sense that for any $\varepsilon > 0$, if $\delta_* = \varepsilon + \Delta \cdot \min\{\alpha - 1, \beta - 1\}$ then we can choose Δ so that policy (α, β) may reject an input of length L where $|L - 1| < \delta_*$. Finally, it is easy to see that if $|L - 1| > \Delta \cdot \max\{1 + \alpha, 1 + \beta\}$ then policy (α, β) will surely reject. Hence the choice of δ^* is valid. Again, it is the best possible. **Q.E.D.**

For instance, policy $(1/2, 1/2)$ is a $(0, 3\Delta/2)$-policy, and policy $(3/2, 2)$ is a $(\Delta/2, 3\Delta)$-policy.

POLICY FOR FIXED δ, Δ. Given δ, Δ, we want to choose α, β such that policy (α, β) is a (δ_*, δ^*)-policy that is compatible with equation (1) and $\delta^* - \delta_*$ is minimized. If we use the bounds (2) from the above lemma, then we see that $\delta^* - \delta_* \geq 2\Delta$, with equality iff $\alpha = \beta$. When $\alpha = \beta$, "policy (α, β)" will be simply called "policy α". Clearly we can choose a compatible (δ_*, δ^*)-policy with $\alpha = \beta$. For compatibility, we can choose any $\alpha > 1$ such that $(\alpha - 1)\Delta < \delta < (\alpha + 1)\Delta$. However, we can choose δ_*, δ^* that are sharper than (2) for every fixed δ, Δ. The following is a "natural choice" for α, β:

$$\alpha_0 = \left\lfloor \frac{1+\delta}{\Delta} \right\rfloor - \frac{1}{\Delta}, \quad \beta_0 = \frac{1}{\Delta} - \left\lceil \frac{1-\delta}{\Delta} \right\rceil. \qquad (3)$$

If we set $m_0 := \lceil \frac{1-\delta}{\Delta} \rceil$ and $m_1 := \lfloor \frac{1+\delta}{\Delta} \rfloor$, then policy (α_0, β_0) amounts to "accept iff $n_0 \in [m_0, m_1]$". Let us assume that $\Delta \leq \delta$. Then $m_0 < m_1$. We may verify that

$$1 - \delta \in [(m_0 - 1)\Delta, m_0\Delta], \quad 1 + \delta \in [m_1\Delta, (m_1 + 1)\Delta].$$

Thus, for this is a (δ_*, δ^*)-policy for some $0 < \delta_* < \delta^* < \infty$. Indeed, $\delta_* = \min\{1 - m_0\Delta, m_1\Delta - 1\}$ and $\delta^* = \max\{1 - (m_0 - 1) \cdot \Delta, (m_1 + 1) \cdot \Delta - 1\}$. Moreover,

$$\Delta \leq \delta^* - \delta_* < 2\Delta.$$

PENALTY MODELS. To distinguish among the possible explicit policies, we may introduce a penalty model. Consider a simple linear model parametrized by the real numbers $p_A, p_R \geq 0$: if we wrongly accept then we pay a penalty of $p_A \cdot \max\{L - 1 - \delta, 1 - \delta - L\}$, and if we wrongly reject then we pay a penalty of $p_R \cdot \min\{1 + \delta - L, L - 1 + \delta\}$. See figure 4.

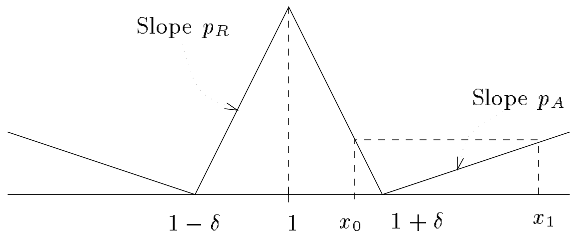

Figure 4: *Linear Penalty Model*

We define a policy to be **optimal** if it minimizes the maximum penalty over all decisions. In general, we want to choose a (δ_*, δ^*)-policy in which the maximum penalty for wrong acceptance equals the maximum penalty for right acceptance. In case Δ is Egyptian, this amounts to the choice of $\delta_* = x_0$ and $\delta^* = x_1$ such that $(\delta - x_0) \cdot p_R = (x_1 - \delta) \cdot p_A$, as illustrated in figure 4. This is not enough to pin down x_0, x_1. But using a formula such as (2), the additional constraint $\delta^* - \delta_* = 2\Delta$ will pin down x_0, x_1 uniquely. Of course, such x_0, x_1 may not exist in case the optimum policy is not δ_*-neoliberal for any $\delta_* > 0$.

Note that the "natural" policy (3) may not be optimal if p_A, p_R are very different. However, it is easy to see that the optimal policy is only small modification of the natural policy. That is, the optimal policy is to accept iff $n_0 \in [m_0', m_1']$ for some m_0', m_1' satisfying $|m_i' - m_i| \leq 1$ $(i = 0, 1)$.

4 Statistical Analysis

Evaluation of policies is a staple subject of mathematical statistics (policies are usually called "tests" or "decision rules" in that setting). To deploy statistical methods, we must first introduce probabilistic assumptions. Indeed, there is a very natural one in our context: note that the probe value n_0 is[4] either $\lfloor L/\Delta \rfloor$ or $\lceil L/\Delta \rceil$. Assuming that every position of B on the Δ-grid is equally likely, then the probability of getting the value $\lfloor L/\Delta \rfloor$ is

$$p := \frac{L}{\Delta} \mathbf{mod}\, 1$$

where $x \,\mathbf{mod}\, 1$ denotes the fractional part of a real number x. Let us now generalize our decision problem by performing m probes, for some given $m \geq 1$. Let the average value of the m probes be the random variable X_m. Clearly, $\lfloor L/\Delta \rfloor \leq X_m \leq \lceil L/\Delta \rceil$. We note a slight problem in case X_m is an integral value n_0: this happens iff all our m probe values are equal to a single value n_0, and we have no way of telling whether this n_0 is $\lfloor L/\Delta \rfloor$ or $\lceil L/\Delta \rceil$. To get around this, we will assume that we do know the value of $\lfloor L/\Delta \rfloor$ through some prior experiment, and let X denote the "shifted" random variable $X_m - \lfloor L/\Delta \rfloor$. Thus

$$X \in \{0, 1/m, \ldots, (m-1)/m, 1\}$$

and the probability $\mathtt{Pr}\{X = i/m\}$ $(i = 0, \ldots, m)$ has the binomial distribution $B(m, p)$.

Returning to the classification problem (1), we may now think of this as the problem of testing the **null hypothesis** $H : L \in [1 \pm \delta]$ against its **complementary hypothesis** $K : L \notin [1 \pm \delta]$. The standard tool here is the Neyman-Pearson theory [5, 3] for testing a pair (H, K) of hypothesis. Let us now describe this setting. We need a randomized test (or policy)

$$\phi : X \mapsto [0, 1]$$

where $\phi(X)$ corresponds to the probability of rejecting the hypothesis. Let X have probability distribution F_θ where $\theta \in \Theta \subseteq \mathbb{R}$ is an unknown real parameter. Suppose H (resp., K) correspond to the hypothesis that $\theta \in [\theta_0, \theta_1]$ (resp., $\theta \notin [\theta_0, \theta_1]$). A **type I error** occurs when we wrongly reject H and a **type II error** occurs when we wrongly accept H. For any $0 < \alpha < 1$, we call ϕ a **most powerful test** (MPT) of **significance** α for (H, K) if the probability of type II error is minimized by ϕ, subject to the probability of **type I error** being at most α. The Neyman-Pearson theory studies conditions for the existence of MPT's.

Let us now assume that $\{F_\theta : \theta \in \Theta\}$ is a "one-parameter exponential family" of distributions, *i.e.*, F_θ has the form

$$C(\theta) e^{Q(\theta)T(x)} h(x)$$

where $Q(\theta)$ is strictly increasing. A generalization of the Neyman-Pearson lemma [5, theorem 6, p.101] says that for each α, there is a MPT of significance α for the pair (K, H) (not the pair (H, K)).

We may apply this generalized lemma for our problem where $\Theta = [0, 1]$ and F_θ is the binomial distribution $B(m, \theta)$. (The reader may verify that this is a one-parameter exponential family.) For any chosen significance α, this lemma tells us how to construct a corresponding MPT. We note some shortcomings in this approach: (i) There does not seem to be a convenient algorithm for constructing the MPT ϕ. (ii) Although we freely choose the significance α, we have no *a priori* bound on the type II error. (iii) Finally, this lemma only applies when K is the null hypothesis, but not when H is the null hypothesis.

5 Multiprobe Strategies

In the preceding study, we focused on the policy π since the sampling σ and algorithm α are trivial. We now make the algorithm non-trivial by allowing "multi-probes". For any $k \geq 0$, a **k-multiprobe** is a sequence of $k + 1$ probes in which:

(a) The initial (zeroth) probe is, as before, an arbitrary placement of B on the Δ-grid.

(b) Each of the subsequent k probes are obtained by user-determined "shifts" of B. The shifts can depend on the results of previous shifts.

[4] If L/Δ is integer, we must interpret $\lceil L/\Delta \rceil$ to mean $1 + \lfloor L/\Delta \rfloor$. We shall employ this interpretation throughout this discussion.

A "shift" is determined by a real number σ where $0 < |\sigma| < \Delta$ and amounts to translating B by the distance σ. The **outcome** of a probe is, instead of being just an integer, the set of grid points that are covered by B. The choice of σ_i ($i = 1, \ldots, k$) in the ith shift may be chosen to depend on the outcomes of the previous i probes.

E.g., a 1-multiprobe allows us to shift B once after its initial placement on the grid. Intuitively, it is clear that the best choice for this shift is $\sigma = \pm\Delta/2$.

Note that there is an uncertainty of Δ at each end of the segment B after the initial (zeroth) probe. The total uncertainty is therefore 2Δ. With each additional probe, we reduce the uncertainty at both ends of the segment B. For any $k \geq 0$, let U_k denote the minimax total uncertainty (in units of Δ) at the two ends of B after k additional probes.

Clearly $U_0 = 2$. For $k = 1$, with $\sigma = 1/2$ for the shift (which is clearly optimal), we achieve the bound $U_1 = 1$. For $k = 2$, the choice $\sigma = 1/3$ for the first shift leads to the bound $U_2 \leq 2/3$. It is also not hard to see that this is optimal. In general, let $0 < \sigma_k \leq 1/2$ denote the first shift in an optimal k-multiprobe. The key remark is this: *once the first shift in a k-multiprobe is given, the optimal strategy for the remaining shifts can essentially be solved by induction on k.* The initial values of σ_k and U_k can be calculated using this remark; they are illustrated in the following table.

k	1	2	3	4	5	6	7	8
σ_k	$\frac{1}{2}$	$\frac{1}{3}$	$\frac{5}{17}$	$\frac{15}{49}$	$\frac{79}{275}$	$\frac{237}{787}$	$\frac{1261}{4409}$	$\frac{3783}{12601}$
U_k	1	$\frac{2}{3}$	$\frac{8}{17}$	$\frac{16}{49}$	$\frac{64}{275}$	$\frac{128}{787}$	$\frac{512}{4409}$	$\frac{1024}{12601}$

The asymptotic behavior of U_k is not immediate from these initial values. However, we may note that

$$2^{1-k} \leq U_k \leq 2^{1-\lfloor k/2 \rfloor}.$$

The first inequality follows from the fact that in the best case, each probe may simultaneously reduce both uncertainty intervals in half. The second inequality follows from the fact that if we treat the two uncertainty intervals independently, then we need two probes to reduce both intervals by half. In a separate paper [16], we determine the asymptotic behavior of U_k.

6 Final Remarks

1. The case study here is intended to raise issues that should be of greater practical interest in 2 or 3 dimensions. E.g., the probe model can clearly be generalized to higher dimensions and applied to other classification problems.

2. Our probe model is "RISC-like" (following Canny and Goldberg). It has the useful property that there is no corresponding "localization problem". In contrast, many probe methods require that certain **localization problems** be solved before the measurement proper can begin. E.g., the method may require that the body be first put in some restricted placement, or it may require estimating the Euclidean transformation from the actual placement of the object to be measure to some special placement.

3. We have looked at deterministic policies, except when we discussed statistical approaches. One expects non-deterministic policies to be more powerful in general. It may also be possible to introduce "uncertainty" into the model, independent of any probabilistic assumptions. We explored a possible statistical approach and this seems to leave interesting unresolved questions.

References

[1] Shaw C. Feng and Theodore H. Hopp. A review of current geometric tolerancing theories and inspection data analysis algorithms. Technical Report NISTIR-4509, National Institute of Standards and Technology, U.S. Department of Commerce. Factory Automation Systems Division, Gaithersburg, MD 20899, February 1991.

[2] Theodore H. Hopp. Computational metrology. In *Proc. 1993 International Forum on Dimensional Tolerancing and Metrology*, pages 207–217, New York, NY, 1993. The American Society of Mechanical Engineers. CRTD-Vol.27.

[3] Maurice G. Kendall and Alan Stuart. *The advanced theory of statistics: volume 2, inference*

and relationship. Hafner Publishing Company, New York, second edition, 1973.

[4] David P. Leech and Albert N. Link. The economic impacts of NIST's software error compensation research. NIST Planning Report 96-2, National Institute of Standards and Technology, TASC, Inc., 1101 Wilson Boulevard, Suite 1500, Arlington, Virginia 22209, June, 1996.

[5] E.L. Lehmann. *Testing statistical hypotheses.* John Wiley and Sons, New York, second edition, 1986.

[6] James D. Meadows. *Geometric Dimensioning and Tolerancing.* Marcel Dekker, Inc, 1995.

[7] K. Mehlhorn, T.C. Shermer, and C.K. Yap. Probing for near-centers and approximate roundness, to appear, 1996. Preliminary version was presented at the ASME Workshop on Tolerancing & Metrology, University of North Carolina at Charlotte, June 21-23, 1995. See URL `http://cs.nyu.edu/cs/faculty/yap`.

[8] Alvin G. Neumann. The new ASME Y14.5M standard on dimensioning and tolerancing. *Manufacturing Review*, 7(1):16–23, March 1994.

[9] Aristides A. G. Requicha. Toward a theory of geometric tolerances. *International Journal of Robotics Research*, 2(4):45–60, 1983.

[10] U. Roy, C.R. Liu, and T.C. Woo. Review of dimensioning and tolerancing: representation and processing. *Computer-aided Design*, 23(7):466–483, 1991.

[11] U. Roy and X. Zhang. Establishment of a pair of concentric circles with the minimum radial separation for assessing rounding error. *Computer Aided Design*, 24(3):161–168, 1992.

[12] Vijay Srinivasan and Herbert B. Voelcker, editors. *Dimensional Tolerancing and Metrology*, 345 East 47th Street, New York, NY 10017, 1993. The American Society of Mechanical Engineers. CRTD-Vol. 27.

[13] Richard K. Walker and Vijay Srinivasan. Creation and evolution of the ASME Y14.5.1M standard. *Manufacturing Review*, 7(1):16–23, March 1994.

[14] ANSI Y14.5M-1982. *Dimensioning and tolerancing.* American Society of Mechanical Engineers, New York, NY, 1982.

[15] Chee K. Yap. Exact computational geometry and tolerancing metrology. In David Avis and Jit Bose, editors, *Snapshots of Computational and Discrete Geometry, Vol.3.* McGill School of Comp.Sci, Tech.Rep. No.SOCS-94.50, 1994. A Volume Dedicated to Godfried Toussaint.

[16] Chee K. Yap and Ee-Chien Chang. A simultaneous searching problem, August, 1996. pickup site, URL `http://cs.nyu.edu/cs/faculty/yap`.