Struct Multidisc Optim (2015) 52:269–279
DOI 10.1007/s00158-015-1237-9

RESEARCH PAPER

L-dominance: An approximate-domination mechanism for adaptive resolution of Pareto frontiers

B. J. Hancock[1] · T. B. Nysetvold[1] · C. A. Mattson[1]

Received: 28 May 2014 / Revised: 30 January 2015 / Accepted: 7 March 2015 / Published online: 13 April 2015

Abstract In Evolutionary Multi-objective Optimization (EMO), the mechanism of ϵ-dominance has received significant attention because of its ability to guarantee convergence near the Pareto frontier and maintain diversity among solutions at a reasonable computational cost. A noticeable weakness of this mechanism is its inability to vary the resolution it provides of the Pareto frontier based on the frontier's tradeoff properties. We therefore propose a new mechanism—L-dominance, based on the Lamé curve—as an alternative to ϵ-dominance in EMO. The geometry of the Lamé curve naturally supports a greater concentration of Pareto solutions in regions of significant tradeoff between objectives. This variable resolution of solutions allows an algorithm using L-dominance to generate fewer solutions to describe the Pareto frontier as a whole while maintaining a desired concentration of solutions where the frontier requires greater detail. The L-dominance mechanism is analyzed theoretically and by simulation on five test problems, and is shown to result in increasingly significant computational gains as the dimensionality of problems increases.

Keywords Multi-objective optimization · Evolutionary optimization · Approximate dominance · ϵ-dominance · Lamé curve · Smart Pareto set

✉ C. A. Mattson
mattson@byu.edu

1 Department of Mechanical Engineering, Brigham Young University, Provo, UT 84602, USA

1 Introduction

Nearly all practical design problems include multiple objectives. Often, these objectives are in conflict with each other. In recent years, one of the most popular approaches for performing numerical optimization under such conditions has been Evolutionary Multi-objective Optimization (EMO). EMO methods facilitate the simultaneous evolution of multiple potential solutions toward the Pareto frontier—the region of the design space that represents all nondominated solutions. An introduction to EMO methods and principles can be found in Deb (2008) and Coello et al. (2007). The goals of these methods are typically to (a) converge upon the Pareto-optimal frontier, (b) maintain diversity among the converged upon solutions, and (c) achieve the first two goals at a reasonable computational cost (Deb et al. 2003).

To achieve these goals, various mechanisms have been proposed, implemented, and analyzed. These mechanisms are not standalone methods for performing multi-objective optimization, but rather series of calculations that can be incorporated into other complete optimization algorithms in order to assist in accomplishing one or more of the three aforementioned desirable goals. One category of such optimization mechanisms is *approximate domination* mechanisms, which modify the traditional definition of domination in some way such that convergence, diversity, or efficiency is improved.

The ϵ-dominance mechanism, first introduced by Laumanns et al. (2002), has received a significant amount of attention in the discipline in recent years because of its ability to achieve all three of these goals when implemented in conjunction with an EMO algorithm. When ϵ-dominance is being used, the design space is divided into boxes with dimensions equal to ϵ, hereafter referred to as ϵ-boxes. No more than one solution is allowed per ϵ-box, and no solutions are allowed in any ϵ-box which is dominated by another occupied ϵ-box. (The details of this mechanism are further discussed in Section 2.2). Limiting the number of solutions in each region ensures diversity, and convergence is guaranteed once all of the boxes that contain the Pareto frontier have been occupied. The calculations associated with this mechanism are computationally benign, and so the third EMO goal is achievable as well. In multiple studies, the ϵ-based versions of EMO algorithms have been shown to outperform their non-ϵ counterparts in terms of improved efficiency, distribution, and rate of convergence (Kollat and Reed 2006; Hadka and Reed 2012; Kollat and Reed 2005).

Despite these apparent strengths, one very significant weakness of ϵ-dominance is its apathy concerning the tradeoff properties of the Pareto frontier. Multiple studies have shown that decision-makers tend to select solutions from regions of the Pareto frontier where the tradeoff ratio between two objectives is changing most quickly—in other words, where the Pareto frontier "bulges" in some direction (Bechikh et al. 2010; Das 1999; Branke et al. 2004). For example, three regions of significant tradeoff have been identified in Fig. 1. These regions are often described as *knees* of the Pareto frontier or *compromise regions*. However, because these terms have been defined many different ways in the literature (Schutze and Laumanns 2008; Das 1999; Rachmawati and Srinivasan 2009; Deb 2003), to avoid ambiguity we will refer to these regions simply as "regions of significant tradeoff." For the decision-maker to adequately understand the full range of possibilities available to him or her, greater resolution (i.e., higher concentration of solutions, and therefore a smaller ϵ-box size) is required in these regions so that the detailed geometry of the Pareto frontier can be sufficiently represented. But by definition, ϵ-boxes are the same size throughout the design space, and using this same maximum resolution over the entire frontier can result in great computational inefficiencies, particularly in high-dimensional spaces.

In this paper, we propose the application of a new mechanism that provides many of the benefits of ϵ-dominance while also allowing for variable resolution of the Pareto frontier in order to better capture its geometry in regions of significant tradeoff. This mechanism employs a shape known as a Lamé curve that was introduced and utilized in another application of multi-objective optimization in a previous publication by the authors (Hancock and Mattson 2013). By preserving diversity through the use of Lamé curves instead of ϵ-boxes, the mechanism of L-dominance achieves the three primary goals of EMO methods while naturally promoting greater resolution of the Pareto frontier in regions of significant tradeoff.

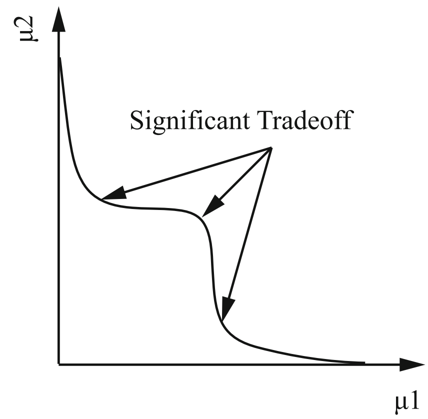

Fig. 1 Three regions of significant tradeoff between objectives are identified on a Pareto frontier in two dimensions

The remainder of this paper is organized as follows: We begin in Section 2 by briefly reviewing the concepts of ϵ-dominance and variable Pareto frontier resolution. In Section 3 we introduce the mechanism of L-dominance, describe how it can be applied, and discuss its advantages and disadvantages. In Section 4 we compare the computational efficiency of using L-dominance vs. ϵ-dominance. In Section 5 we discuss the significance of the results and make suggestions for future research. Finally, in Section 6 we offer concluding remarks.

2 Technical preliminaries

This section reviews the theory, definition, and proposed variations of ϵ-dominance, which serves as the foundation for the L-dominance mechanism. In Section 2.1 we define three categories of domination. In Section 2.2 we provide the definition of ϵ-dominance as it was originally introduced. Then in Section 2.3 we briefly discuss three ϵ-dominance variations and one other mechanism that have already been proposed in the literature as means of providing varied resolution of the Pareto frontier.

2.1 Types of dominance

For the definitions in this section, let $X, Y \in \mathbb{R}^n$, and assume all objectives are to be minimized.

2.1.1 Strict dominance

Point X is said to *strictly* dominate point Y if X is strictly better than Y in all n objectives:

$$\forall i \in \{1, ..., n\} : X_i < Y_i \tag{1}$$

2.1.2 Weak dominance

Point X is said to *weakly* dominate point Y if X is better than or equal to Y in all n objectives, and strictly better in at least 1 objective:

$$\forall i \in \{1, ..., n\} : X_i \leq Y_i \tag{2}$$

$$\exists j \in \{1, ..., n\} : X_j < Y_j \tag{3}$$

2.1.3 Approximate dominance

Point X is said to *approximately* dominate point Y if Y is located within the defined region of approximate equality surrounding X, and/or X weakly dominates Y. The region of approximate equality is described by an ϵ-box for ϵ-dominance, and a Lamé curve for L-dominance. These shapes are defined in Section 2.2 and Section 3.1, respectively. As a note, any shape used to define a region of approximate equality must have the symmetric property that if a solution A is located in the region of approximate equality for solution B, then solution B must also be located in the region of approximate equality for solution A. Consequently, each approximate domination mechanism necessarily has a policy for selecting the more desirable solution when two solutions would otherwise mutually approximately dominate each other.

2.2 ϵ-dominance

When using ϵ-dominance, the user provides a value for ϵ that represents the minimum amount of change in an objective that he or she considers to be significant. Then using intervals of size ϵ, the design space is partitioned into a grid of ϵ-boxes. Each solution is contained within one of these ϵ-boxes, and the coordinates of that box become the values by which weak-dominance is assessed. Consider, for example, the design space presented in Fig. 2, wherein all objectives are to be minimized and $\epsilon = 10$. While solution C does not weakly dominate solution B—not all Cartesian coordinates of C (15, 12) are less than or equal to the Cartesian coordinates of B (13, 27)—C *does* ϵ-dominate B, since the ϵ-box coordinates of C (2, 2) do weakly dominate the ϵ-box coordinates of B (2, 3). Thus, in this design space, only solutions A, C, and E are ϵ-nondominated. When two solutions are discovered in the same ϵ-box (such as solutions C and D), the tie-breaker policy is that the one with the smaller Euclidean distance to the optimal corner (the lower right corner in Fig. 2) approximately dominates the other.

When using this approach, anchor point solutions near the ends of the Pareto frontier (such as solution G in Fig. 2) are often approximately dominated and removed. In most situations, however, it is desirable to include these solutions in the final set so that the full shape of the Pareto frontier is represented, (even though resolution may vary throughout the frontier). To achieve this end, a variant on traditional ϵ-domination may be used wherein strict dominance of ϵ-boxes is required instead of weak dominance when removing other solutions. Under these conditions, solutions A, B, C, E, and G would be ϵ-nondominated. In the remainder of this paper, all references to ϵ-dominance may be understood to refer to the strict dominance version of the mechanism.

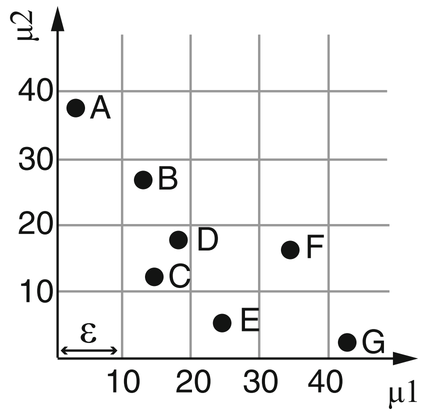

Fig. 2 In this design space where $\epsilon = 10$, solution C does not weakly dominate solution B, but *does* ϵ-dominate it. In this design space, only solutions A, C, E are ϵ-nondominated

Once again, ϵ-dominance is referred to as a mechanism rather than a method because the mere existence of ϵ-boxes is not sufficient to generate a Pareto frontier. Rather, calculations of ϵ-dominance are used within existing EMO methods when determining which solutions to archive or remove from generation to generation. This flexibility has allowed ϵ-dominance to be incorporated into a wide variety of methods. A few of these include ϵ-MOEA (Deb et al. 2003), ϵ-NSGA-II (Kollat and Reed 2005), ϵ-MyDE (Santana-Quintero and Coello 2005), and Adaptive-ϵ Box (Kowatari et al. 2012).

2.3 ϵ-dominance variations

As was mentioned in Section 1, a recognized weakness of ϵ-dominance is its inability to adjust the resolution of generated solutions within the design space. It maintains diversity equally in all regions of the Pareto frontier, regardless of where signficant tradeoff is or isn't occurring. Consequently, a number of ϵ-dominance variations have been proposed in the years since its inception.

The original creators of ϵ-dominance recommended using a different value of ϵ_i for each objective i in the design space. This potentially changes the ϵ-boxes from squares to rectangles, but still clearly lacks the ability to increase the resolution of solutions only where it is needed.

Jin and Wong (2003) proposed using a grid of dynamically adjusted ϵ-boxes (see Fig. 3a). This approach, called Adaptive Rectangle Archiving (ARA), allows for the grid lines that define ϵ-boxes to move along the objective axes in order to adjust the density of solutions in the design space. However, according to Laumanns et al. (2002), the original creators of ϵ-dominance, "for ensuring the convergence property it is important not to move or translate any of the box limits; in other words, the assignment of the elements to the boxes must stay the same." Thus, the improvement in solution density that can come from allowing the limits of ϵ-boxes to shift during the course of the optimization comes at the cost of one of the three primary objectives of EMOs—a guarantee of convergence.

Five years after the introduction of ϵ-dominance, Hernández-Díaz et al. (2007) introduced the concept of *Pareto-adaptive ϵ-dominance*, wherein an approximation of the Pareto frontier is obtained and associated with an equation of the form:

$$x_1^p + x_2^p + \ldots + x_n^p = 1 : \quad 0 < p < \infty \tag{4}$$

The calculated value for p is then used along with user-provided values for the desired number of solutions and the "speed of variation" to compute values of ϵ_i that vary throughout the space according to a geometric sequence. This approach results in smaller ϵ-boxes where a single region of significant tradeoff of the Pareto frontier is projected to be (see Fig. 3b). However, this method requires prior knowledge of the shape of the Pareto frontier (i.e., function calls spent on creating an approximation), assumes a symmetric geometry, and cannot be adjusted to identify multiple regions of significant tradeoff.

In addition to these ϵ-dominance variations created for obtaining variable Pareto frontier resolution, the use of a post-processing mechanism called a *smart Pareto filter* was proposed by Mattson et al. (2004). This mechanism steps through a densely populated set of Pareto solutions and systematically removes all solutions that lie in the *Practically Insignifcant Tradeoff* (PIT) region of an already accepted solution. The dimensions of the PIT region are chosen by the user specifically for the problem at hand. After the filter has been applied, the remaining set of solutions has the desirable property of having greater resoultion in regions of significant tradeoff and is called a *smart Pareto set*. Nevertheless, this mechanism can only be applied after all candidate solutions have been found, and it requires a method capable of first generating a high concentration of solutions in all desired regions of the Pareto frontier.

Clearly, the ability of a method or mechanism to identify and exploit regions of significant tradeoff is very desirable within the discipline. In the remainder of this paper, we introduce and analyze the L-dominance mechanism, which has this ability while maintaining many of the properties that have made ϵ-dominance a frequently used mechanism within evolutionary multi-objective optimization.

3 The L-dominance mechanism

In this section, we first introduce the mathematical definition of the Lamé curve, then demonstrate how it can be used to enable L-dominance calculations, and finally discuss the advantages and disadvantages of L-dominance compared to ϵ-dominance. The fundamental concept behind L-dominance is the replacement of ϵ-boxes with a shape that naturally lends itself to a greater concentration of solutions in regions of significant tradeoff. This occurs due to the geometry of the Lamé curve, which has the smallest thickness in directions of high tradeoff (e.g., 45° in a uniformly scaled space), and the largest thickness in directions of low tradeoff (e.g., the vertical and horizontal directions) over the specified domain of $p = [0, 2]$. Using Lamé curves

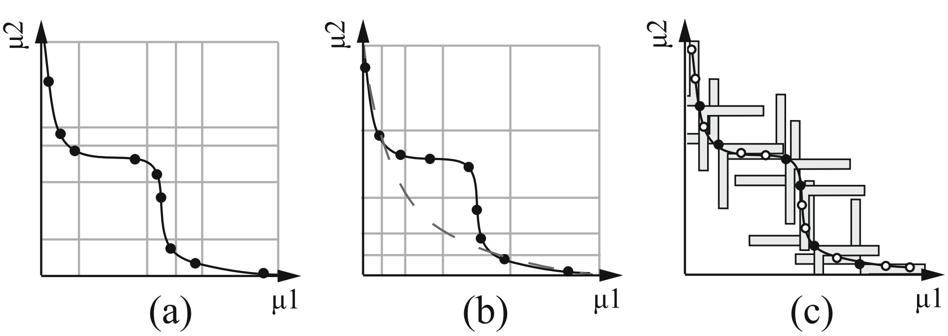

Fig. 3 Pareto sets resulting from (**a**) Adaptive Rectangle Archiving with dynamically adjusted ϵ-boxes, (**b**) Pareto-adaptive ϵ-dominance with varied ϵ-box sizes based on an approximation of the Pareto frontier, and (**c**) the use of a smart Pareto filter, which removes all solutions (empty circles) that lie in the PIT region (shaded gray) of an accepted solution (filled circles)

to define approximate dominance rather than ϵ-boxes, diversity and convergence are still guaranteed, but exploitation of regions of significant tradeoff is also enabled.

3.1 The mathematical definition of the Lamé curve

The Lamé curve (also sometimes called a superellipse) is defined by the equation

$$\left(\sum_{i=1}^{n} |A_{i,i}(x_i - d_i)|^p\right)^{\frac{1}{p}} = 1 \qquad (n \geq 2) \tag{5}$$

where

$$\mathbf{A} = \begin{bmatrix} \frac{1}{a_1} & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \frac{1}{a_n} \end{bmatrix} \tag{6}$$

and $\mathbf{d}$ is a vector from the origin to the center point of the Lamé curve. Algebraically, this represents the p-norm (Rynne 2007) of $\mathbf{x}$, offset by $\mathbf{d}$ (based on its location in the design space), multiplied by a transformation matrix $\mathbf{A}$ (for scaling purposes), and set equal to 1. When used as an EMO mechanism, the user-defined variables a_i and p of the Lamé curve allow the user to determine the distribution of the Pareto solutions that will be generated for that particular problem. Each value a_i (sometimes called the semi-diameter of a Lamé curve) corresponds to objective i in the problem and may be interpreted as the amount of change in that objective that would constitute a significant difference between two solutions in the user's mind if all other objectives remain practically unchanged (see Fig. 4). Accordingly, larger values for a_i will result in larger approximate domination boundaries, and therefore fewer solutions in the final Pareto set.

The parameter p affects the curvature of the approximate domination boundaries and therefore controls the extent to which significant tradeoff between objectives is required in order for two solution nearby each other to both remain in the Pareto set. The effect of p on the shape of the approximately dominated region is illustrated in Fig. 5. While the method will work for any value of p between 0 and 2, it is assumed that for most purposes, the user will select a value between 0 and 1. At $p = 0$, the area of the approximately dominated region approaches 0. At $p = 2$, the Lamé curve becomes an ellipse (or higher dimensional version of an ellipse), resulting in an approximately even distribution of solutions over the entire Pareto frontier.

The computational cost of generating a Lamé curve or of identifying whether a solution lies within or without it is constant. However, as is made evident by Fig. 5, smaller values of p will result in smaller approximately dominated regions (just as smaller values of ϵ result in smaller ϵ-boxes). Smaller approximately dominated regions with either mechanism will result in more solutions being found on the Pareto frontier, and consequently more function calls being made overall in the optimization routine. Thus, as would be expected, the choice of values for the user-defined parameters in either ϵ-dominance or L-dominance will affect the overall quantity of solutions in the final set and therefore the total computational cost of the routine. However, while lowering the value of p will require greater computational cost overall, it will also result in a lower computational cost than what would be required if the same resolution were being sought with ϵ-boxes. This result is discussed further in Sections 4 and 5.

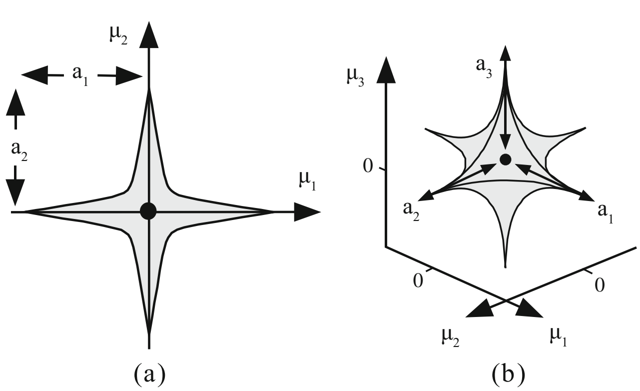

Fig. 4 The Lamé curves for a design space in a design space with (**a**) 2 objectives and (**b**) 3 objectives

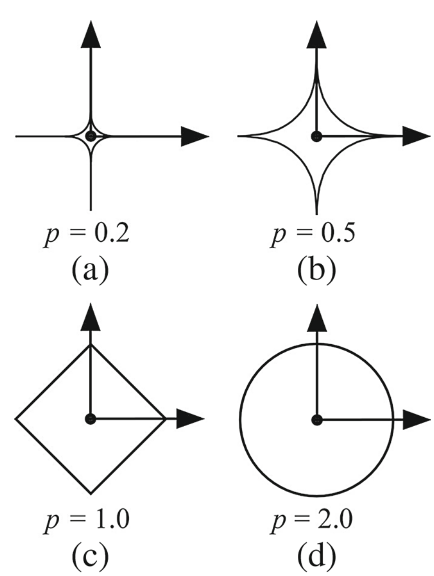

Fig. 5 The effect of the user-defined parameter p on the shape of the approximately dominated region bounded by a Lamé curve

3.2 How to apply the L-dominance mechanism

Calculations for L-dominance occur in the same place where any ϵ-dominance calculations occur within an ϵ-based method, and the output for either mechanism is the same—a Boolean value that indicates whether a candidate solution is approximately dominated or not, and should therefore be added to an archive, replace an existing solution, or be dropped, depending on the mode of operation of the algorithm. This means that L-dominance can easily be incorporated into the wide variety of methods that currently utilize an approximate domination mechanism.

As described in the definition of approximate dominance given in Section 2.1, a new solution is L-dominated if it is either located within the Lamé curve of an existing Pareto solution or weakly dominated. Like ϵ-boxes, Lamé curves remain stationary in the design space once they are created, thereby preserving a guarantee of convergence (since deterioration of the set of Pareto solutions is not allowed to occur). Unlike ϵ-boxes, however, Lamé curves have no other coordinates in the space besides the Cartesian coordinates of the center point. Thus, domination is evaluated using Cartesian coordinates rather than ϵ-box coordinates. In situations where two solutions lie within each other's approximately dominated region and neither weakly dominates the other, the tie-breaker policy for L-dominance is to simply accept the solution that was discovered first.

3.3 Comparison of L-dominance and ϵ-dominance

The L-dominance and ϵ-dominance mechanisms have much in common in the way that they satisfy the three objectives of EMO. First, both guarantee convergence near the Pareto frontier. According to Laumanns et al. (2002), in EMO methods with archiving strategies, convergence is assured if deterioration of the archive cannot occur. Deterioration occurs when elements of the current archive are dominated by solutions that were in the archive at some previous time. Because the bounds of approximately dominated regions remain fixed for both mechanisms and no new solutions may be added to the archive unless they are nondominated by the current archive, convergence is consequently guaranteed. Once the Pareto frontier is entirely contained within either Lamé curves or ϵ-boxes, no new solutions may be added. This stagnation of the Pareto set that corresponds with no new solutions being found can easily be used as a termination condition for an optimization algorithm. For further discussion of convergence properties of EMO algorithms, see Rudolph and Agapie (2000).

Second, both mechanisms guarantee diversity—based on the user-defined values for a_i and p (for L-dominance) or ϵ_i (for ϵ-dominance), a certain minimum spacing between nearby solutions is required for all to remain in the Pareto set.

And third, the calculations for both mechanisms are computationally benign. Because the Lamé is describable with a single equation, the determination of whether or not a new solution resides in an existing Lamé curve is a very simple check to perform.

Aside from satisfying these three objectives of EMOs, L-dominance provides a number of additional benefits beyond ϵ-dominance and its proposed variations. As has already been mentioned, L-dominance allows a greater concentration of solutions to occur in significant tradeoff regions of the Pareto frontier. And, as will be shown in Section 4, the benefits of this property in terms of computational gains grow significantly as the number of dimensions in the problem increases. There is no limit to how many of these significant tradeoff regions can be exploited (since variable resolution comes naturally as a result of the shape of the Lamé curve), and the mechanism requires no prior information about the shape of the Pareto frontier to work. Unlike ϵ-dominance, where adjacent solutions can be located immediately on either side of the line between ϵ-boxes, there is a minimum distance that adjacent solutions are guaranteed to have between them in L-dominance. Finally, because weak domination is checked using Cartesian coordinates rather than ϵ-box coordinates, anchor point solutions—which generally occur where at least one coordinate is changing very slowly—are more likely to be discovered when using L-dominance.

There are also some potential disadvantages that come from using L-dominance instead of ϵ-dominance. While L-dominance provides the user with more flexibility by allowing him or her to specify the curvature of the Lamé curve, the proper selection of this parameter is less immediately intuitive than the selection of the lengths of an ϵ-box. This may be viewed as either an additional cognitive burden being placed on the user or as an opportunity for the user to include more of his or her experience and knowledge of the problem in the optimization formulation. If variable resolution is desired, then any value for p in the allowable range will produce a shape more naturally inclined to promote this goal than an ϵ-box, since the direction of significant tradeoff is now the thinnest part of the approximately dominated region instead of the thickest. And if the user would rather not have variable resolution or not have to make a decision on what value of p to use, then setting $p = 2$ will yield Lamé curves in the shape of a circle (or its higher dimensional corollary), which will yield an equal distribution of solutions just as ϵ-dominance would.

The other potential disadvantage of L-dominance is that when two solutions are discovered in the same ϵ-box, the tie-breaker method may be employed to select the more preferred of the two, since the location of the ϵ-box remains

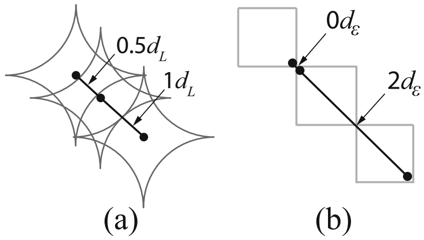

Fig. 6 The average distance between solutions in regions of high tradeoff is shown to range between [0.5, 1] diagonal lengths for Lamé curves, and between [0, 2] diagonal lengths for ϵ-boxes

unchanged either way. Because Lamé curves are centered on individual solutions as they appear rather than being defined at the beginning of the optimization like ϵ-boxes, the first solution in a given Lamé curve must be the winner of all "ties" between two solutions within the same curve that do not weakly dominate each other, since the boundaries of an approximate equality region must remain fixed to guarantee convergence.

4 Simulations

4.1 Simulation setup

One way to assess the benefits of using L-dominance instead of ϵ-dominance is to compare the resulting Pareto sets when these two mechanisms are utilized within the same EMO algorithm, on the same problem, and with similar settings. For each of the simulation problems, we first performed an optimization using L-dominance. Given the shape of the Lamé curves used, we were able to calculate the average distance between solutions in a region of significant tradeoff. As shown in Fig. 6, adjacent solutions in a converged Pareto frontier can vary between 0.5 and 1 times the "diagonal" of a Lamé curve. Given the symmetry of the Lamé curve, this yields an average distance of 0.75 times the diagonal (Knuth 1986), yielding the equation:

$$\bar{d_L} = (0.75)(2)(n^{-\frac{1}{p}})\sqrt{\sum_{i=1}^{n} \mathrm{a}_i^2} \tag{7}$$

When using ϵ-boxes, on the other hand, adjacent solutions can vary between approximately 0 and 2 times the diagonal of an ϵ-box. This results in an average distance of:

$$\bar{d_\epsilon} = \sqrt{n}(\epsilon) \tag{8}$$

where n is the number of dimensions in the problem. Setting the two average distances between solutions equal to each other, we then solved for ϵ and used that value to perform an optimization with ϵ-dominance. Thus, although the stochastic nature of evolutionary algorithms resulted in slightly different resolution in the Pareto frontier with each execution, the average distance between solutions in regions of high tradeoff was the same for both mechanisms. By observing differences in the required number of function calls to achieve the same average resolution, the computational benefits of using L-dominance vs. ϵ-dominance can be quantified.

For the simulations described in this paper, the approximate domination mechanisms were used within the ϵ-MOEA algorithm developed by Deb et al. (2003). This algorithm was chosen both for its relatively simple structure and because it was one of the first to utilize ϵ-dominance and is consequently well-known in the field. Because the focus of this paper is on the approximate domination mechanisms rather than the optimization algorithms within which they are used, we refer those interested in the details of the ϵ-MOEA algorithm to Appendix A or the original publication.

A sample simulation is represented graphically in Fig. 7, which shows an illustration of the Pareto frontier of a simple two-dimensional design space with boundaries in the shape of a circle. First, L-dominance is used to generate a Pareto set such as the one shown in Fig. 7a. The average distance between solutions in the region with maximum significant tradeoff is recorded as $\bar{d}$, as indicated. Then, using the formulas given above, the necessary value for ϵ is calculated and used to generate a Pareto set with the same $\bar{d}$ using ϵ-dominance, as shown in Fig. 7b. As indicated by the plots, the average distance between solutions in the region of interest is nearly identical. However, while the Pareto set generated by L-dominance is able to reduce resolution in the relatively flat regions of the Pareto frontier (as seen from the increasing magnitude of $\bar{d}$ in the upper left and lower right corners of Fig. 7a), ϵ-dominance experiences approximately uniform distribution everywhere.

4.2 Problem descriptions and tabulated results

The three multi-objective problems that were used to compare ϵ-dominance to L-dominance are summarized in Table 1. Each problem was solved in two, three, and four dimensions. Between these three problems, Pareto frontiers that are convex, concave, and disjoint are all represented. The problem identified as Circle represents an n-dimensional design space bounded by a circle (or its higher dimensional corollary, such as a sphere or hypersphere), which is centered on the point that has all units set equal to 1 (such as (1,1) in two dimensions). The DTLZ2 and DTLZ7 problems were introduced by Deb et al. (2002) as scalable test problems for benchmarking MOEAs. The DTLZ2 has a concave Pareto frontier, and DTLZ7 has 2^{M-1} disjoint regions that collectively make up its Pareto frontier. Both problems are described in greater detail in the aforementioned paper.

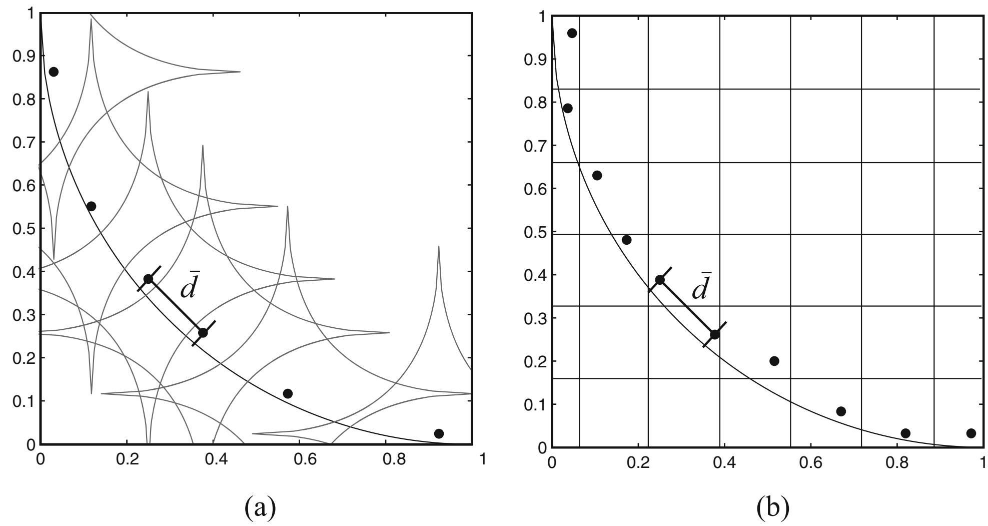

Fig. 7 An illustration of Pareto sets that could be generated by (**a**) L-dominance and (**b**) ϵ-dominance using parameter values designed to yield the same average resolution $\bar{d}$ in the region of maximum significant tradeoff on the Pareto frontier of the Circle problem in 2D

The simulation results are summarized in Table 2. The first section gives the values for the arbitrary user-defined parameter a_i for L-dominance and the corresponding value of ϵ that would result in the same average resolution in high tradeoff regions. A value of $p = 0.5$ was used for all L-dominance simulations. The second section reports the total number of solutions generated to reach convergence, and the final section reports the total number of function calls performed for each mechanism. All reported quantities of solutions and function calls are the average of 25 simulations. Because we are primarily interested in observing general trends, this level of precision was sufficient for this study.

5 Discussion of simulation results

The most significant trend to recognize from the simulation results is shown in bold in Table 2, and that is that as the number of dimensions increases, so does the reduction in the number of function calls required to reach convergence for L-dominance compared to ϵ-dominance. This is to be expected, since the number of ϵ-boxes in a design space grows exponentially as dimensionality increases. In four dimensions, the ratio of function calls performed is as low as 16 %, and engineering problems with many more than four objectives are not uncommon.

We note that the ratio of the number of solutions for each mechanism is not directly proportional to the ratio of the number of function calls required. This suggests that the cost (in function calls) of producing a single solution can vary by problem due to the geometry of the Pareto frontier and how that geometry interfaces with the shape of approximate domination that is being utilized. For example, unlike ϵ-dominance, which has predefined regions of approximate equality, L-dominance creates its regions of approximate equality as new solutions are found. This can

Table 1 This table provides the governing equations for the three test problems that were used to conduct simulations

Name	Problem	Geometry
Circle	$f_{m=1:M} = x_m$ subject to $h(\mathbf{x}) \geq 0$ where $h(\mathbf{x}) = \sum_{i=1}^{M}(x_i - 1)^2$	Convex
DTLZ2	$f_1 = (1+g)\prod_{i=1}^{M-1}\cos(x_i\pi/2)$ $f_{m=2:M-1} = (1+g)\left(\prod_{i=1}^{M-m}\cos(x_i\pi/2)\right)\sin(x_{M-m+1}\pi/2)$ $f_M = (1+g)\sin(x_i\pi/2)$ $g = \sum_{x_i\in\mathbf{x}_M}(x_i - 0.5)^2$	Concave
DTLZ7	$f_{m=1:M-1} = x_m$ $f_M = (1+g)\left(M - \sum_{i=1}^{M-1}\left[\frac{f_i}{1+g}(1+\sin(3\pi f_i))\right]\right)$ $g = 1 + \frac{9}{\lvert\mathbf{x}_M\rvert}\sum_{x_i\in\mathbf{x}_M} x_i$	Disjoint

Each problem is scalable to any number of dimensions, represented here as M. The bounds on all independent variables are [0, 1]

Table 2 This table compares the number of solutions and function calls required to reach convergence for each mechanism on each of three test problems in 2D, 3D, and 4D

	2D Circle	DTLZ2	DTLZ7	**3D** Circle	DTLZ2	DTLZ7	**4D** Circle	DTLZ2	DTLZ7
Parameters									
a_i	0.25	0.5	0.5	1	1.5	1.5	2	4	4
ϵ_i	0.094	0.188	0.188	0.167	0.250	0.250	0.188	0.375	0.375
$\bar{d}$	0.133	0.266	0.266	0.289	0.433	0.433	0.375	0.750	0.750
# Solutions									
L-dominance	9	8	6	15	17	13	22	18	8
ϵ-dominance	15	14	9	63	49	21	278	177	15
# Function Calls									
L-dominance	670	408	801	552	1,143	818	593	430	584
ϵ-dominance	1,006	711	1,572	1,282	2,955	2,272	3,471	2,738	3,554
Ratio	**0.67**	**0.57**	**0.51**	**0.43**	**0.39**	**0.36**	**0.17**	**0.16**	**0.16**

User-defined parameters were set such that both mechanisms produce Pareto sets with the same average resolution in regions of high tradeoff. Of particular interest is the decreasing ratio of the number of function calls for L-dominance compared to ϵ-dominance as dimensionality increases

result in numerous small regions between proximate Lamé curves where new solutions can be discovered, thereby providing more frequent opportunities for new members to be recognized as approximately nondominated even if the amount of progress toward the true Pareto frontier is the same for both mechanisms. Consequently, if the convergence criterion for an algorithm is simply based on the rate of change in the archive of nondominated solutions, this may effect a difference in the perceived rate of progress between the two mechanisms, and therefore the time required for convergence. The convergence criterion used for these simulations was, indeed, simply a maximum allowable number of function calls (300) without the discovery of a new Pareto solution. We therefore hypothesize that the use of more sophisticated convergence criteria could result in even greater gains for L-dominance than what is already shown in Table 2.

While not reflected in Table 2, where $p = 0.5$ for all L-dominance simulations, we also noticed that generally, as p decreased (resulting in "skinnier" Lamé curves), the relative advantage of L-dominance over ϵ-dominance in terms of function calls increased. This was true only up to some minimum bound (usually between 0.25 and 0.33) where reductions in cost leveled out. We hypothesize once again that this may be the result of the design space becoming more "fragmented" as tighter Lamé curves are used, causing a greater cost per Pareto solution due to more frequent resetting of the stagnancy counter that corresponds to the convergence criterion. We suggest investigation of this phenomenon as additional future work.

Differences in performance between L-dominance and ϵ-dominance depend to some extent on the algorithm within which they are used (e.g. ϵ-MOEA, ϵ-NSGAII, etc.), as well as the setup and parameters of the optimization (e.g., convergence criteria, constraint handling, mutation rate, etc.). These factors will obviously affect the rate of convergence as well as the average number of function calls per solution. In these simulations, because the focus was simply to compare two mechanisms under standard operating conditions, very little was done to optimize the overarching algorithm that applied the approximate domination mechanisms (in this case, ϵ-MOEA). As with most optimization tasks, for best results, the user will identify the algorithm and optimization parameters most suitable for the given problem. Thus, while results will vary for each setup and problem to which these mechanisms are applied, based on the theory laid out in this paper we would expect the general trends found by our simulations and shown in Table 2 to hold.

Finally, we propose future work with connections to a number of developments in EMO that have been published in recent years. Emmerich (2007) have proposed a gradient-based/evolutionary hybrid method, which attempts to maximize the hypervolume being dominated by a set of solutions. While that method does not utilize approximate domination calculations, it is feasible that the measure of dominated hypervolume could be adjusted/weighted to reflect its position in the design space relative to tradeoff properties. Similarly, Zhang and Li (2007) and Giagkiozis et al. (2013) have proposed decomposition methods, which divide the multi-objective optimization problem into a set of neighboring subproblems. While both were presented with the assumption that a uniform distribution is desired, the definition of a well-distributed Pareto set could also be modified to reflect the desirability of a variable resolution on the Pareto frontier, using the equation of the Lamé curve as a guide for where greater resolution is required.

6 Conclusion

The mechanism of ϵ-dominance provides many benefits in EMO, including guaranteed convergence and diversity preservation at a reasonable computational cost. However, without the ability to adjust the resolution of the Pareto frontier according to its tradeoff properties, the methods employing ϵ-dominance often experience unnecessary inefficiencies. In this paper, we introduced a new mechanism, L-dominance, which maintains the positive qualities of ϵ-dominance while naturally increasing the concentration of solutions in regions of significant tradeoff. This is accomplished primarily through the use of Lamé curves instead of ϵ-boxes as the region of approximate equality for an approximate domination mechanism. The value of the proposed mechanism has been shown theoretically and by simulation results. In particular, increasingly significant reductions in computational cost were seen as the dimensionality of the problems increased. Based on these results, we confidently recommend the use of L-dominance in place of ϵ-dominance as a diversity/convergence mechanism where variable resolution is desired, especially in high-dimensional EMO problems.

Acknowledgments We would like to recognize the National Science Foundation (Grant CMMI-1301247) for funding this research.

Appendix A: The ϵ-MOEA Algorithm

The ϵ-MOEA algorithm was first introduced by Deb et al. (2003). In this algorithm, as with many EMO algorithms, there are two groups of solutions being maintained: the general population $P(t)$ and the archive population $A(t)$, where t indicates the generation number. An initial population (often randomly generated) fills $P(0)$, and the ϵ-nondominated solutions from the population become the initial archive population $A(0)$. One member from each population is then chosen for mating. The solution from $A(t)$, called a, is chosen at random. To choose a solution from P(t), called p, two randomly selected solutions are checked for dominance. If one dominates the other, then that solution is chosen; otherwise, the choice between the two is random. Solutions a and p are then mated to create λ offspring solutions. In this paper and the original ϵ-MOEA paper, $\lambda = 1$. These offspring candidate solutions, called c_i for $i = 1...\lambda$, are then compared to both populations for possible inclusion.

For inclusion in the general population, a standard dominance check is performed between c_i and the current general population. If c_i is dominated by any member of $P(t)$, then it is not accepted. If c_i dominates ones or more members, then it replaces one of them at random. If c_i is non-dominated with respect to $P(t)$, then it replaces a population member chosen at random. In all cases, the size of $P(t)$ remains constant.

For inclusion in the archive population, an ϵ-dominance check is performed between c_i and the current archive population $A(t)$. If c_i is ϵ-dominated by any member of $A(t)$, then it is not accepted. If c_i ϵ-dominates one or more members, then it replaces one of them at random. So far, this is nearly identical to the inclusion process for the general population. In the case that c_i is ϵ-nondominated with respect to $A(t)$, however, one of two actions may be taken. If c_i occupies a previously unoccupied ϵ-box, it is accepted into the archive. If c_i occupies an ϵ-box that is currently occupied by an archive solution, then a tie breaker is performed to determine which will remain and which will be removed. If one dominates the other in the usual sense, the nondominated solution remains. If the two solutions are nondominated with respect to each other, then the accepted solution is the one that has the smaller Euclidean distance to the corner of the ϵ-box that is in the preferred direction for all objectives (e.g., the bottom left corner in a two-dimensional space where both objectives are being minimized).

This procedure is repeated for a specified number of generations and the final archive population is returned as the Pareto set of solutions. As described above, this algorithm is a steady-state, elitist MOEA. The size of $P(t)$ is constant through all generations, and while $A(t)$ has the potential to grow in size over the course of the algorithm, its size is also bounded by the fact that the Pareto frontier must pass through a finite number of ϵ-boxes, and each of these may contain only a single archive member. Furthermore, because no solution can be removed from the archive until a solution that either dominates it or ϵ-dominates it has been added in its place, deterioration of the Pareto set is prevented at all times.

References

Bechikh S, Said L, Ghedira K (2010) Searching for knee regions in multi-objective optimization using mobile reference points. In: Proceedings of the 2010 ACM symposium on applied computing

Branke J, Deb K, Dierolf H, Osswald M (2004) Finding knees in multi-objective optimization. In: Parallel Problem Solving from Nature-PPSN VIII. Springer, pp 722–731

Coello C, Lamont G, Van Veldhuizen D (2007) Evolutionary algorithms for solving multi-objective problems. Springer

Das I (1999) On characterizing the “knee” of the pareto curve based on normal-boundary intersection. Structural Optimization 18(2-3):107–115

Deb K (2003) Advances in evolutionary computing. Springer, Berlin Heidelberg

Deb K (2008) Introduction to evolutionary multiobjective optimization. In: Multiobjective optimization. Springer, pp 59–96

Deb K, Thiele L, Laumanns M, Zitzler E (2002) Scalable multi-objective optimization test problems. Proceedings of the 2002 IEEE Congress on Evolutionary Computation

Deb K, Mohan M, Mishra S (2003) A fast multi-objective evolutionary algorithm for finding well-spread pareto-optimal solutions. KanGAL report 2003002

Emmerich M (2007) Graident-based/evolutionary relay hybrid for computing Pareto front aapproximation maximizing the S-metric. Springer, Berlin Heidelberg

Giagkiozis I, Purshouse R, Fleming P (2013) Generalized decomposition. Springer, Berlin

Hadka D, Reed P (2012) Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. Evol Comput 20(3):423–452

Hancock B, Mattson C (2013) The smart normal constraint method for directly generating a smart pareto set. Struct Multidiscip Optim 48:763–775

Hernández-Díaz A, Santana-Quintero L, Coello CC, Molina J (2007) Pareto-adaptive ε-dominance. Evol Comput 15(4):493–517

Jin H, Wong M (2003) Adaptive diversity maintenance and convergence guarantee in multiobjective evolutionary algorithms. In: The 2003 Congress on Evolutionary Computation CEC’03, vol 4. IEEE, pp 2498–2505

Knuth D (1986) Metafont. Addison-Wesley

Kollat J, Reed P (2005) The value of online adaptive search: a performance comparison of nsgaii, ε-nsgaii and ε-moea. In: Evolutionary multi-criterion optimization. Springer, pp 386–398

Kollat J, Reed P (2006) Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design. Adv Water Resour 29(6):792–807

Kowatari N, Oyama A, Aguirre H, Tanaka K (2012) A study on large population moea using adaptive ε-box dominance and neighborhood recombination for many-objective optimization. In: Learning and intelligent optimization. Springer, pp 86–100

Laumanns M, Thiele L, Deb K, Zitzler E (2002) Combining convergence and diversity in evolutionary multiobjective optimization. Evol Comput 10(3):263–282

Mattson C, Mullur A, Messac A (2004) Smart pareto filter: obtaining a minimal representation of multiobjective design space. Eng Optim 36:721–740

Rachmawati L, Srinivasan D (2009) Multiobjective evolutionary algorithm with controllable focus on the knees of the pareto front. IEEE Trans Evol Comput 13:810–824

Rudolph G, Agapie A (2000) Scalable multi-objective optimization test problems. In: Zalzala A, Eberhart R (eds) Congress on Evolutionary Computation, vol 2. IEEE Press, pp 1010–1016

Rynne B (2007) Linear functional analysis. Springer

Santana-Quintero L, Coello CC (2005) An algorithm based on differential evolution for multi-objective problems. Int J Comput Intell Res 1(1):151–169

Schutze O, Laumanns M (2008) Approximating the knee of an MOP with stochastic search algorithms. Springer-Verlag

Zhang Q, Li H (2007) Moea/d: a multiobjective evolutionary algorithm based on decomposition. IEEE Trans Evol Comput 11(6):712–731