cost network-flow problems has been investigated in Fruhwirth, Burkard, and Rote [6].

Whereas in convex programming a locally good approximation in the neighborhood of the optimal solution is sufficient, a global approximation with a uniform error bound on the whole range is required in the other applications mentioned above: In bicriteria and parametric programming one is interested in the optimal value function (or efficient point curve, respectively) as a whole.

Several authors have considered the problem of approximating convex functions $h(t)$ by a piecewise linear function. If h is twice continuously differentiable ($h \in C^2$), Phillips [13] finds approximations with a given error bound on the whole domain of h. Phillip's procedure requires the computation of roots of nonlinear equations to find the points at which the piecewise approximation and the original function coincide. (These points will subsequently be called knots). By using this procedure iteratively one can also find the best approximation with respect to a given number of knots. Cox [4] showed that Phillip's results hold if h is just continuously differentiable ($h \in C^1$), but that the computations of the knots can be simplified if $h \in C^2$. Thakur's [20–22] objective is to find approximations such that the difference between the optimal objective values of the original mathematical program and its approximate version is bounded by a given error. He solves a series of smaller piecewise linear problems and uses an equidistant partitioning of the domain of h to develop convergent procedures. Salem and Elmaghraby [17] develop local approximation methods for convex, decreasing, but nonquadratic functions, and apply their results to project networks. Further references for approximation by piecewise linear functions include Kao and Meyer [10], Sonnevend [18, 19], and the textbooks of Bazaraa and Shetty [2] and Rockafellar [14]. Geoffrion [7] develops a more general theory of approximating objective functions in mathematical programming.

The geometric problem of approximating a plane convex figure by a convex polygon, which is essentially the same as our problem, has also received great attention in the literature, both from a geometric theoretical view point (cf. the overview of Gruber [8]) as well as in the applied areas of image processing and computational geometry. A nice overview of algorithms from this area is given in Kurozumi and Davis [11], cf. also Rote [16].

In this article we develop the sandwich algorithm for computing a global approximation for a convex function h in a given interval $[a, b]$. At any stage of the algorithm piecewise-linear lower and upper approximations $l(t)$ and $u(t)$, respectively, of $h(t)$ are known on the whole domain $[a, b]$ of h and will satisfy

$$l(t) \le h(t) \le u(t), \qquad \text{for all } a \le t \le b.$$

The algorithm terminates when the global error is less than a prespecified ε; i.e.,

$$\sup_{a \le t \le b} \{u(t) - l(t)\} < \varepsilon.$$

Alternatively, the sandwich algorithm may be terminated if a given number of knots have been computed. In this case one will be able to compute a priori an achievable error bound of the approximation.