FFT Acceleration of the Rectangular Dock Problem *

Tom Korsmeyer †

The Continuous Problem

The hydrodynamic behavior of offshore structures large enough to be used for floating airports or military bases is difficult to model computationally because of their size relative to typical ocean wave lengths. Boundary-element methods (BEM) are usually applied to these problems, but their computational cost (cpu time expended and memory allocated) is high. Conventional BEM rely on the set-up and solution of dense linear systems so their cost is at least order N^2, in which N is the number of unknowns. The techniques used to mitigate the analysis cost of these structures have been:

- high-order BEM, which is more efficient than low-order [4];
- low-order BEM with *ad hoc* assumptions to reduce the number of unknowns (see [2]);
- hierarchical clustering of interactions between similar parts of the structure with BEM to compute the representative interactions (see [2]).

The first of these approaches retains the order-N^2 cost of conventional BEM, but can achieve higher accuracy for a given N than a low-order BEM. The latter two of these approaches can reduce the cost to less than order N^2, but at the expense of either provable accuracy or general application.

The key to solving these dense matrix problems with reduced computational cost is to approximate them with sparse systems, thereby reducing both the set up and the solution cost. An algorithm for solving BEM formulations that effectively sparsifies the linear system while providing *a priori* error bounds is the precorrected-FFT method [5]. In this approach, the interaction of nearby elements is computed directly, and that of the remainder is approximated by efficient calculations on a uniform grid. The efficiency of the grid-based interaction relies on the fact it may be cast as a circulant matrix-vector product. A circulant matrix represents convolution and so circulant-matrix computations may be done with reduced cost by the FFT.

We consider a structure for a floating airport that consists of multiple rectangular barges such that the entire structure has length and beam considerably greater than the draft and the spacing between adjacent barges is negligible. This structure may be analyzed hydrodynamically or hydroelasticly as an infinitely thin plate on the free surface (the "dock" problem). This particular floating-airport structure leads naturally to the circulant form when analyzed by a BEM and may be solved extremely efficiently by a method that exploits the low cost of the FFT. While admittedly a special case, the presentation of this solution demonstrates the ideas behind the more general approach of precorrected-FFT.

The problem is linearized so that the total potential is

$$\Phi(x) = \Phi_I(x) + \Phi_S(x) + \sum_j \Phi_j(x), \tag{1}$$

in which $\Phi_I(x)$ is an incident wave potential, $\Phi_S(x)$ is the scattering potential, and $\Phi_j(x)$ are the radiation potentials for rigid and non-rigid modes. These potentials all satisfy similar boundary-value and integral formulations, so we consider, for example, the diffraction problem (for $\Phi_S(x)$ $\;x : x \in \mathbf{R}^3, x_3 \le 0$) for a plate in the plane $x_3 = 0$ of length L and width B in a fluid of depth H. Let S_M denote the plate surface $-\frac{L}{2} \le x_1 \le \frac{L}{2}$, $-\frac{B}{2} \le x_2 \le \frac{B}{2}$, and $x_3 = 0$. We will solve this problem in the frequency domain, that is, we will work with the complex amplitude $\phi_S(x)$ defined in

$$\Phi_S(x,t) = \mathrm{Re}\{\phi_S(x)e^{i\omega t}\} \tag{2}$$

for a particular radian frequency ω.

*This work is supported by ONR Grant N00014-97-1-0827, under the direction of NFESC; and the Industry Consortium *Numerical Analysis of Wave Loads on Offshore Structures,* sponsored by: Chevron, Exxon, Mobil, Norsk Hydro, NSWC-CD, OTRC, Petrobras, Saga, Shell, Statoil.

†Research Laboratory of Electronics Massachusetts Institute of Technology (xmeyer@mit.edu)

The scattering potential $\phi_S(x)$ satisfies the Laplace equation

$$\nabla^2 \phi_S(x) = 0, \tag{3}$$

and the combined (kinematic and dynamic) boundary conditions on $x_3 = 0$

$$-K\phi_S(x) + \frac{\partial}{\partial z}\phi_S(x) = \frac{i\omega}{\rho g}p(x) \quad x \in S_M, \qquad -K\phi_S(x) + \frac{\partial}{\partial z}\phi_S(x) = 0 \quad x \notin S_M, \tag{4}$$

in which K is the infinite-depth wave number corresponding to the frequency ω, *i.e.* $K = \omega^2/g$. A radiation condition makes the problem well posed.

The problem may be recast as a second-kind Fredholm integral equation on the wetted plate surface to be solved for the complex diffraction pressure directly [7]

$$p(x) + \frac{K}{2\pi}\int_{-\frac{L}{2}}^{\frac{L}{2}} d\xi_1 \int_{-\frac{B}{2}}^{\frac{B}{2}} d\xi_2 \, p(\xi) G(\xi; x, \omega) = -\rho g \zeta_I(x, \omega) \qquad x \in S_M, \tag{5}$$

in which $G(\xi; x, \omega)$ is the frequency-domain, free-surface, Green function and $\zeta_I(x, \omega)$ is the incident wave elevation.

The Discrete Problem

For this case of a rectangular plate, if (5) is discretized by a structured mesh of $N = N_L \times N_B$ identical rectangular elements, upon which the pressure is assumed to be constant, and the equation is enforced at the element centroids, then the matrix of the resulting linear system is doubly-Toeplitz in form. More precisely, in

$$Tp = \zeta \tag{6}$$

($T \in \mathbf{C}^{N \times N}; p, \zeta \in \mathbf{C}^N$) there are N_L^2 blocks in a Toeplitz pattern and these rank-N_B blocks are themselves of Toeplitz form.

A circulant matrix (or simply "circulant") is a special Toeplitz form with the structure:

$$C = \begin{bmatrix} c_1 & c_n & \dots & c_2 \\ c_2 & c_1 & \dots & \vdots \\ \vdots & \vdots & \dots & c_n \\ c_n & \dots & c_2 & c_1 \end{bmatrix}. \tag{7}$$

The important feature of a circulant is that it may be diagonalized by application of the discrete Fourier transform matrix F_n, by

$$D = F_n C F_n^{-1}$$

for C a circulant and $D = \mathrm{diag}(d_i)$ [6]. Not surprisingly, because of the repetitive structure of C, the elements of D can be calculated simply from the matrix-vector product

$$d = F_n c \tag{8}$$

in which c is the first column of C. This may still appear to entail an unacceptable order-N^2 computational cost, however the use of the FFT reduces it to order-$N \log N$ time and order-N storage.

The way to exploit circulants in a problem with Toeplitz structure is that any Toeplitz matrix of rank N can be embedded in a contrived circulant matrix of rank $2N$. Here, the doubly-Toeplitz structure leads to a doubly-circulant matrix of rank $4N$. We will show that it is a worthwhile strategy to work with a linear system of size $4N$ instead of N, as long as the asymptotic cost is reduced from N^2 to $N \log N$ or N.

Consider solving (6) with a Krylov-subspace iterative method like GMRES. As we iterate to find the solution within a certain tolerance ϵ, forming the residual r^k from the k^{th} trial vector p^k requires the matrix-vector product

$$r^k = \zeta - Tp^k. \tag{9}$$

Provided that the system is sufficiently well conditioned, so that $|r^{k_{\max}}| < \epsilon$ for $k_{\max} << N$ or even $k_{\max}$ independent of N, then the solution has order-N^2 cost because of (9).

We can achieve $|r^{k_{\max}}| < \epsilon$ without forming T, and naturally without explicitly performing the N^2 operations of Tp^k, by exploiting the property of circulants noted above. We do this in two distinct ways: (1) we find a doubly-circulant approximation to T, which we denote $C_P \in \mathbb{C}^{N \times N}$, and use it as a preconditioner, and (2) we embed T in a doubly-circulant matrix, which we denote $C_T \in \mathbb{C}^{4N \times 4N}$, and use it to compute residuals instead of T.

For preconditioners, there are a number of circulants that are reasonable approximations to T. An *ad hoc* choice would be to use the $N_B/2$ central diagonals of the blocks, repeating them as necessary to obtain circulant blocks and extend this idea to the block structure as well. This amounts to emphasizing the half of the elements nearby to any evaluation point, while ignoring the half farther away. A more rational choice is to derive a new set of matrix entries from the old by an algorithm that minimizes a norm. It is straightforward to minimize the Frobenius norm [1]. In this case, no elements are ignored; all are included in weighted sums of near and far influences. Given either of these approaches, instead of working with (6), we solve for y in

$$TC_P^{-1}y = \zeta, \qquad p = C_P^{-1}y. \tag{10}$$

The algorithm (without details of GMRES, particularly the setting of y^k) proceeds as follows:

for $N = N_L \times N_B$ elements {
  compute $4N$ entries of c_T and N entries of c_P /* load circulant columns */
}
set $k = 0$, ϵ; **initialize** $r^k = \zeta, y^k$ /* prepare to iterate with an initial guess */
while $|r^k| > \epsilon$ {   /* iterate as long as tolerance is not met */
  $p^k = \text{fft}^{-1}\left([\text{fft}(c_P)]^T \text{fft}(y^k)\right)$ /* use FFT to apply circulant preconditioner */
  $\hat{p}^k = P(p^k)$ /* P an operator zero padding p^k to nullify $c_{Ti} \notin T$ */
  $\hat{\zeta}^k = \text{fft}^{-1}\left([\text{fft}(c_T)]^T \text{fft}(\hat{p}^k)\right)$ /* apply influence matrix to k^{th} pressure in Fourier space */
  $\zeta^k = P^{-1}(\hat{\zeta}^k)$ /* P^{-1} extracts useful part of $\hat{\zeta}^k$, remaining entries ignored */
  if $(|r^k| = |\zeta - \zeta^k| < \epsilon)$ {**save** p^k; $k_{\max} = k$; **quit**}   /* if tolerance met, we're done */
  else {$k = k + 1$; **choose** y^{k+1}}   /* otherwise set new trial vector and repeat */
}

Results and Discussion

The algorithm has been coded so that preconditioner effectiveness, algorithm accuracy, and computational cost may be examined. Figure 1 shows the effect of the two preconditioners as compared to no preconditioning. For long waves, preconditioning is not very important, though both preconditioners are better than none. However as the waves become shorter than approximately $L/16$, not only does preconditioning become very important, but interestingly the norm-minimizing preconditioner is very effective (nearly independent of λ), while the *ad hoc* preconditioner is actually harmful.

The diffraction pressure for a finite-depth case is shown in Figure 2. These results were computed with 16,384 elements and, through comparison with results using 65,536 elements, appear to be converged to graphical accuracy at least. The computational cost required to obtain such results is shown in Table 1. Again, the effectiveness of the norm-minimizing preconditioner is apparent in its very weak dependence on N. Also, we can see that the cpu time expended is increasing slightly faster than N due to the existence of both order-N and order-$N \log N$ operation requirements in the algorithm, and the memory allocated reflects the strictly order-N storage requirement. If a conventional low-order BEM code were applied to the 16,384-element case, the cost would be between 1 and 2 orders of magnitude greater than with the present method, even if 2 planes of symmetry were exploited. However the more general BEM approach probably could achieve a similar accuracy with fewer elements by using a graded mesh.

The point of this specialized example is to demonstrate the computational efficiencies that can be gained from structured decompositions. Given an arbitrary structure for analysis these efficiencies may still be realized. The idea is that only an order-N subset of the computation of the residual is

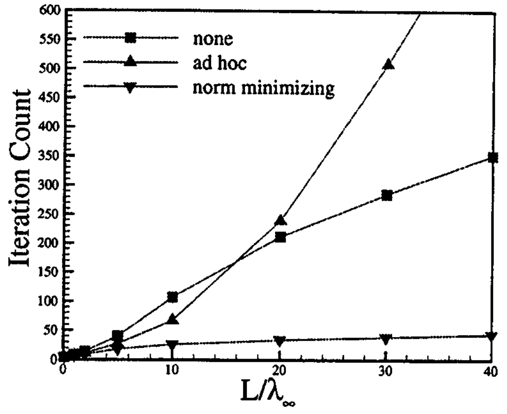

Figure 1: Iterations required ($\epsilon = 0.0005$) for various preconditioning strategies for the diffraction analysis of an $L/B = 5$ plate, in $H/\lambda_\infty = 0.2$ fluid depth and following waves.

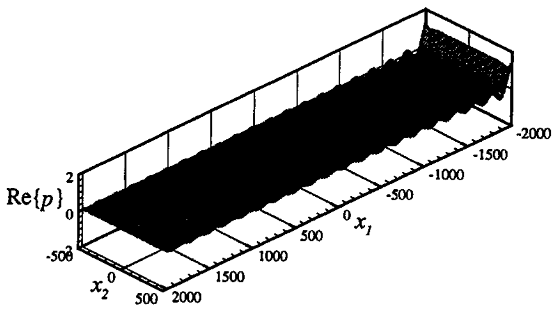

Figure 2: Normalized diffraction pressure on an $L = 4000$m by $B = 1000$m floating airport, in 20m fluid depth and following 200m-length waves. Compare with Figure 4 in [4].

done directly with the boundary elements. The remainder is done on a uniform grid to which the element singularity strengths are projected and from which the potential is interpolated. This method is applicable to integral equations with a variety of kernels, and is demonstrated in [3] for the frequency-domain, free-surface Green function.

N_L	N_B	N	Iterations	CPU Time (sec)	Memory Allocated (Mb)
64	16	1024	33	.6	1
128	32	4096	34	3.1	4
256	64	16,384	37	17.7	16
512	128	65,536	40	91.5	63

Table 1: Computational time *per frequency* and memory allocated (DEC alpha 433MHz, $\epsilon = 0.0005$) for the diffraction analysis of an $L/B = 5$ plate, in $H/\lambda_\infty = 0.2$ fluid depth and following waves. For the 16,384 element case, the estimated time and memory usage for a conventional N^2-cost BEM (exploiting 2 planes of symmetry) is 900 seconds and 340Mb.

References

[1] Tony F. Chan. An optimal circulant preconditioner for Toeplitz systems. *Siam J. Sci. Stat. Comp.*, 9(4), 1988.

[2] M. Kashiwagi, W. Koterayama, and M. Ohkusu. *Hydroelasticity in Marine Technology*, RIAM, Kyushu University, Kyushu, 1998.

[3] T. Korsmeyer, T. Klemas, J. Phillips, and J. White. Fast hydrodynamic analysis of large offshore structures. In *ISOPE'99*, Brest, 1998. To appear.

[4] C.-H. Lee. Wave interaction with huge floating structure. In *8th Int. Conf. on the Behavior of Offshore Structures*, Delft, 1997.

[5] J. R. Phillips and J. K. White. A precorrected-FFT method for electrostatic analysis of complicated 3-D structures. *IEEE Trans. on Computer-Aided Design*, 16(10):1059–1072, 1997.

[6] C. Van Loan. *Computational Frameworks for the Fast Fourier Transform.* SIAM, Philadelphia, 1992.

[7] S. Yamashita. Motions and hydrodynamic pressures of a box-shaped floating structure of shallow draft in regular waves. *J. Soc. of Nav. Arch. of Japan*, 146:165–172, 1977.