

Editorial Manager(tm) for Mathematical Geology
Manuscript Draft

Manuscript Number: MATG188

Title: FFT-based Algorithms for Kriging

Article Type: Original Research

Section/Category:

Keywords: Fast Fourier Transform; Geostatistical Estimation; Spectral Methods; Superfast Toeplitz Solver

Corresponding Author: Dipl.-Ing. Jochen Fritz,

Corresponding Author's Institution: Universität Stuttgart

First Author: Jochen Fritz

Order of Authors: Jochen Fritz; Insa Neuweiler, Dr.rer.nat.; Wolfgang Nowak, Dr.-Ing.

Manuscript Region of Origin:

Abstract: Computational power and storage capacities often pose heavy limitations to the size of the problem able to be addressed in Kriging. For estimation on regular grids and the generalized case of an uncertain and spatially varying mean, we compile a toolbox of FFT-based (spectral) methods for Kriging that is highly efficient in storage and computational complexity. The general theme is to apply, improve or extend existing FFT-based algorithms for basic operations on covariance matrices which apply when covariance matrices have Toeplitz structure. The discussed FFT-based algorithms are easily applicable for the case of regular grids. In case of irregularly scattered data, we trace the problem back to sparse but regular finer grids of measurements. We also present several fast approximations for the estimation variance of Kriged fields that are asymptotically exact for certain limiting cases. The computational efficiency and reduction of storage requirements over existing Kriging algorithms are discussed and demonstrated in test cases.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

FFT-based Algorithms for Kriging*

J.Fritz[†], I. Neuweiler[‡] and W.Nowak[§]

September 11, 2007

Corresponding Author:
Jochen Fritz
Institute of Hydraulic Engineering
Pfaffenwaldring 61
70569 Stuttgart
Germany
jochen.fritz@iws.uni-stuttgart.de
Phone: +49 (0)711 / 685-64667
Fax: +49 (0)711 / 685 - 60430

ABSTRACT

Computational power and storage capacities often pose heavy limitations to the size of the problem able to be addressed in Kriging. For estimation on regular grids and the generalized case of an uncertain and spatially varying mean, we compile a toolbox of FFT-based (spectral) methods for Kriging that is highly efficient in storage and computational complexity. The general theme is to apply, improve or extend existing FFT-based algorithms

*Received ; accepted

[†]Department of Hydromechanics and Modeling of Hydrosystems, Institute of Hydraulic Engineering, Universität Stuttgart, Stuttgart, Germany, Tel. +49 (0)711 685-64667, Jochen.Fritz@iws.uni-stuttgart.de

[‡]Department of Hydromechanics and Modeling of Hydrosystems, Institute of Hydraulic Engineering, Universität Stuttgart, Stuttgart, Germany, Tel. +49 (0)711 685-69165, Insa.Neuweiler@iws.uni-stuttgart.de

[§]Department of Hydromechanics and Modeling of Hydrosystems, Institute of Hydraulic Engineering, Universität Stuttgart, Stuttgart, Germany, Tel. +49 (0)711 685-68218, Wolfgang.Nowak@iws.uni-stuttgart.de

1
2
3 for basic operations on covariance matrices which apply when covariance matrices have
4
5 Toeplitz structure. The discussed FFT-based algorithms are easily applicable for the case
6
7 of regular grids. In case of irregularly scattered data, we trace the problem back to sparse
8
9 but regular finer grids of measurements. We also present several fast approximations for
10
11 the estimation variance of Kriged fields that are asymptotically exact for certain limiting
12
13 cases. The computational efficiency and reduction of storage requirements over existing
14
15 Kriging algorithms are discussed and demonstrated in test cases.
16

17
18 KEY WORDS: Fast Fourier Transform, Geostatistical Estimation, Spectral Methods
19

20 21 22 INTRODUCTION 23

24
25 Spatially distributed quantities such as rainfall intensities, contaminant concentrations or
26
27 hydraulic conductivities are frequently interpolated between scattered measurements by
28
29 Kriging. Especially when considering large data sets of measurements, Kriging can lead to
30
31 systems of equations which are far beyond the storage capacities and computational power
32
33 of contemporary desktop computers. The motivation of this work is to make Kriging fast
34
35 and the required storage capacities low, permitting the solution of very large problems on
36
37 small computers.
38

39
40 In most standard forms of Kriging, there are three computationally most demanding tasks.
41
42 These are (1) to solve an $m \times m$ system of equations that involves the auto-covariance
43
44 matrix of the measurements to obtain the Kriging weights, (2) to perform a superposition
45
46 with the cross-covariance function between measurements and unknowns, weighted by the
47
48 Kriging weights in order to obtain the values at the points of estimation, and (3) to repeat
49
50 this procedure once for the estimator and m times for the estimation variance, where m
51
52 is the number of measurements.
53

54
55 Nearest neighborhood Kriging avoids the solution of large systems of equations by only
56
57 considering measurements within a certain radius around every point of estimation. This
58
59 requires the setting up and solution of as many systems of equations as the number of
60
61 different neighborhoods, and may be very time consuming (e.g.Kitanidis, 1997). Also, an
62
63
64
65

1
2
3 implicit assumption on a moving-average type of mean value has to be made since the
4
5 mean value cannot be estimated globally.
6

7
8 Pegram (2004) published an efficient FFT-based Kriging method. He solves the Krig-
9
10 ing equations using an FFT-based algorithm called Iterative Constrained Deconvolution
11
12 (ICD). It is restricted to the cases of either known or unknown constant mean. For the
13
14 latter case, the sample mean of the measurements is used to estimate the unknown mean
15
16 value of the field. The estimation variance is approximated with only one system of equa-
17
18 tions to be solved, assuming that the measurements are almost uncorrelated. We will
19
20 demonstrate that the ICD algorithm, being formally identical to a non-preconditioned
21
22 steepest descent with an empirically chosen step-size coefficient, can be replaced by a
23
24 more efficient PCG-based algorithm. Also, we will show how to generalize Kriging in the
25
26 FFT-context to the case of uncertain mean, including spatial trends.
27
28

29
30 We exploit the fact that, in most cases, the auto-covariance function of the unknowns is
31
32 assumed to be second-order stationary or at least intrinsic (e.g, Kitanidis 1997), and the
33
34 points of estimation lie on a regular and equispaced grid. This setup leads to covariance
35
36 matrices with symmetric Toeplitz structure (Golub and van Loan 1996). The Toeplitz
37
38 structure can be exploited in terms of storage, because only the first column of the matrix
39
40 has to be stored (Zimmerman 1989). The product of a Toeplitz matrix with a vector is the
41
42 same as discrete convolution of a vector with a corresponding (covariance) function. This
43
44 convolution can be performed quickly using the FFT algorithm for convolution (e.g., van
45
46 Loan 1992). We show how to perform superposition based on this convolution algorithm
47
48 in order to speed up task (2) of evaluating the values at the points of estimation in Kriging.
49
50 With this technique, the associated computational costs can be reduced from the order
51
52 of mn to the order of $n \log_2 n$, where n is the number of estimation points. This is an
53
54 advantage for medium and large numbers of measurements ($m > \log_2 n$).
55
56

57
58 The more fundamental problem, task (1), is to solve the $m \times m$ system of Kriging equations
59
60 with the auto-covariance matrix of the measurements. If the measurements lie on a
61
62 regular grid and the field is second-order stationary, their auto-covariance matrix is again
63
64 a Toeplitz matrix. Per definition of covariance functions, the Toeplitz matrix is symmetric
65

1
2
3 and positive-definite. Solving a Toeplitz system has been the subject of many studies in
4
5 the signal processing community (e.g. Gallivan et al. 1996, Kailath and Sayed 1999,
6
7 Barel et al. 2001). The iterative Toeplitz solver that we find the most promising within
8
9 the Kriging context is the FFT-based Preconditioned Conjugate Gradient (FFT-PCG)
10
11 method with circulant preconditioners. Applying this algorithm has the same prerequisites
12
13 as convolution via FFT, plus positive-definiteness of the involved Toeplitz matrix. We
14
15 refer to Shewchuk (1994) for a graphical review of the PCG algorithm and to Chan and Ng
16
17 (1996) for a comprehensive review on the use of circulant preconditioners therein. Using
18
19 the FFT-PCG algorithm, the Kriging system can be solved with computational effort in
20
21 the order of $m \log_2 m$ instead of m^3 (depending on the conventional solver used).
22
23

24
25 In case the measurements are irregularly spaced, the Toeplitz structure of the auto-
26
27 covariance matrix of measurements is lost. For this case, we present an extension of
28
29 the FFT-PCG algorithm to irregular grids. This extension is similar to the idea used by
30
31 Pegram (2004) in the ICD algorithm, but embedded in a more powerful iterative solver.
32

33
34 Most intrinsic random space functions can be traced back to special cases of second-
35
36 order stationary fluctuations about a spatial varying mean with known base functions
37
38 (i.e., trends) but unknown or uncertain coefficients of the base functions (compare Nowak
39
40 and Cirpka 2004, Nowak and Cirpka 2006, Kitanidis 1993). In order to be as general as
41
42 possible and include such intrinsic cases, we use the generalized form of Kriging for the
43
44 case of a spatially variable and uncertain mean. This case also includes the cases of known
45
46 and unknown mean as limiting cases.
47

48
49 The remaining problem is that task (1) and (2) need to be performed m times in order
50
51 to obtain the estimation variance of Kriging. We show how to reduce the factor of m in
52
53 several levels of trade-off versus accuracy. These approximations are asymptotically exact
54
55 for certain special cases. The fastest options require exactly one additional solution, using
56
57 the approximation for almost uncorrelated measurements as suggested by Pegram (2004),
58
59 or by neglecting boundary effects on large regular grids of measurements. At the end of
60
61 this study, we demonstrate that these methods allow to handle huge Kriging problems
62
63 even on ordinary desktop computers.
64
65

KRIGING WITH UNCERTAIN MEAN

Based on the function-estimate form of Kriging with unknown mean (e.g., Kitanidis 1997 and Kitanidis 1996), Nowak and Cirpka (2004) developed the generalization of Kriging to the case of uncertain mean. Although their original publication dealt with non-linear cokriging-like cases, Kriging is obtained from the described method by using the same covariance function among observed and estimated quantities. The following is a brief summary of the method.

Let \mathbf{s} be an $n \times 1$ multi-Gaussian vector of unknowns (sometimes also called target point values in regression-like problems) with expectation $E[\mathbf{s}] = \mathbf{X}\boldsymbol{\beta}$ and covariance $Cov[\mathbf{s}|\boldsymbol{\beta}] = \mathbf{Q}_{ss}$, where \mathbf{X} denotes the $[n \times p]$ matrix of discrete base functions and $\boldsymbol{\beta}$ the $[p \times 1]$ vector of trend coefficients. For a spatially constant mean of \mathbf{s} , \mathbf{X} is a $n \times 1$ vector with unit entries and $\boldsymbol{\beta}$ is the actual value of the mean. When adding trends, \mathbf{X} is extended by new columns for each additional trend term, with the according coefficient appended to $\boldsymbol{\beta}$. For the case of uncertain mean, the trend coefficients are again set to be Gaussian variables with mean $\boldsymbol{\beta}^*$ and variance $\mathbf{Q}_{\boldsymbol{\beta}\boldsymbol{\beta}}$. For given values of \mathbf{s} , the distribution of $\boldsymbol{\beta}$ is again Gaussian with conditional mean $E[\boldsymbol{\beta}|\mathbf{s}] = \hat{\boldsymbol{\beta}}$ and conditional covariance $\mathbf{Q}_{\boldsymbol{\beta}\boldsymbol{\beta}|\mathbf{s}}$.

Consider further \mathbf{Y} an $m \times 1$ vector of measurements (sometimes also called control point values in regression-like problems) which are sampled from \mathbf{s} according to

$$\mathbf{Y} = \mathbf{H}\mathbf{s} + \mathbf{r}. \quad (1)$$

Here, \mathbf{H} is a $m \times n$ sampling defined as

$$\mathbf{H}_{i,j} = \begin{cases} 1 & \text{for } \mathbf{x}_i = \mathbf{x}_j \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

with \mathbf{x}_i the coordinates of the i -th measurement location and \mathbf{x}_j the coordinates of the j -th estimation point, and \mathbf{r} is an $m \times 1$ vector of random measurement errors. The

1
2
3 measurement errors have zero mean and $m \times m$ covariance matrix \mathbf{R} , typically a scalar
4
5 matrix (i.e., a diagonal matrix with constant values along the diagonal). Interpreting
6
7 the sampling matrix \mathbf{H} as a sensitivity matrix in linear error propagation, the following
8
9 identities hold:

$$10 \quad \mathbf{Q}_{ys} = \mathbf{H}\mathbf{Q}_{ss} \quad (3)$$

$$11 \quad \mathbf{Q}_{sy} = \mathbf{Q}_{ss}\mathbf{H}^T \quad (4)$$

$$12 \quad \mathbf{Q}_{yy} = \mathbf{H}\mathbf{Q}_{ss}\mathbf{H}^T + \mathbf{R}, \quad (5)$$

13
14
15 where \mathbf{Q}_{sy} is the $n \times m$ cross-covariance matrix between the unknowns \mathbf{s} and measurements
16
17 \mathbf{Y} , and \mathbf{Q}_{yy} is the $m \times m$ auto-covariance matrix of the measurements. In this notation,
18
19 the Kriging estimate $\hat{\mathbf{s}}$ is given by

$$20 \quad \hat{\mathbf{s}} = \begin{bmatrix} \mathbf{Q}_{ys} \\ \mathbf{X}^T \end{bmatrix}^T \begin{bmatrix} \boldsymbol{\xi} \\ \hat{\boldsymbol{\beta}} \end{bmatrix}. \quad (6)$$

21
22 The $m \times 1$ vector of Kriging weights $\boldsymbol{\xi}$ (also called reciprocal data by Pegram 2004) and
23
24 the $p \times 1$ vector of trend coefficients $\hat{\boldsymbol{\beta}}$ are taken from the solution of the Kriging system

$$25 \quad \begin{bmatrix} \mathbf{Q}_{yy} & \mathbf{H}\mathbf{X} \\ (\mathbf{H}\mathbf{X})^T & -\mathbf{Q}_{\beta\beta}^{-1} \end{bmatrix} \begin{bmatrix} \boldsymbol{\xi} \\ \hat{\boldsymbol{\beta}} \end{bmatrix} = \begin{bmatrix} \mathbf{Y} \\ -\mathbf{Q}_{\beta\beta}^{-1}\boldsymbol{\beta}^* \end{bmatrix}. \quad (7)$$

26
27 The associated estimation variance $\hat{\boldsymbol{\sigma}}$ is the $n \times 1$ vector on the diagonal of the conditional
28
29 covariance matrix

$$30 \quad \mathbf{Q}_{ss|y} = \mathbf{Q}_{ss} - \begin{bmatrix} \mathbf{Q}_{ys} \\ \mathbf{X}^T \end{bmatrix}^T \begin{bmatrix} \mathbf{Q}_{yy} & \mathbf{H}\mathbf{X} \\ (\mathbf{H}\mathbf{X})^T & -\mathbf{Q}_{\beta\beta}^{-1} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{Q}_{ys} \\ \mathbf{X}^T \end{bmatrix}. \quad (8)$$

31
32 The special case of unknown mean is recovered by setting $\mathbf{Q}_{\beta\beta}^{-1} = 0$ and using a vector of
33
34 unit entries for \mathbf{X} . The case of known mean is included by setting $p = 0$, i.e., by omitting
35
36 all rows and columns that refer to $\boldsymbol{\beta}$.

1
2
3 The covariance of \mathbf{s} for uncertain β is $\mathbf{G}_{\mathbf{ss}} = \mathbf{Q}_{\mathbf{ss}} + \mathbf{X}\mathbf{Q}_{\beta\beta}\mathbf{X}^T$, where $\mathbf{G}_{\mathbf{ss}}$ is a generalized
4 covariance matrix $\mathbf{G}_{\mathbf{ss}}$ (Kitanidis 1993). If $\mathbf{Q}_{\mathbf{ss}}$ is second-order stationary, then $\mathbf{G}_{\mathbf{ss}}$ is at
5 least intrinsic. By adequate choice of \mathbf{X} , the most common intrinsic cases can be expressed
6 via generalized covariance matrices, which in turn can be decomposed into second-order
7 stationary parts plus terms related to an uncertain mean value.
8
9

10
11 The expensive tasks referred to the introduction are (1) solving the Kriging system (Eq. 7)
12 to obtain the Kriging weights, (2) perform the superposition $\mathbf{Q}_{\mathbf{sy}}\boldsymbol{\xi}$ in Eq. (6) to evaluate
13 the estimate, and (3) evaluate $m + p$ equivalents to task one and two to obtain the
14 estimation variance (Eq. 8) .
15
16
17
18
19
20
21
22

23 **TOOLBOX OF SPECTRAL METHODS**

24
25
26
27 In this section, we provide and extend a collection of FFT-based methods which speed
28 up matrix operations for so-called Toeplitz matrices. The connection between spatial
29 estimation and Toeplitz matrices is explained in the following section. The basic trick of
30 all FFT-based methods is that their periodic counterparts, called circulant matrices, have
31 highly exploitable properties in the Fourier space. Their most central property is that
32 the discrete Fourier transform of their first column equals their eigenvalues and that the
33 eigenvectors of all circulant matrices are combined in the discrete Fourier matrix (Varga
34 1954). This so-called diagonalization theorem has also been proved by Trapp (1973).
35 Hence, Toeplitz matrices are first converted to circulant matrices in a step called periodic
36 embedding, then FFT-techniques are applied, followed by conversion back to the Toeplitz-
37 equivalent in a step called extraction. Illustrated reviews of this procedure are provided
38 by Kozintsev (1999) or by Nowak (2005).
39
40
41
42
43
44
45
46
47
48
49
50
51
52

53 **Exploitable structure of covariance matrices**

54
55
56
57 For finely resolved Kriging problems in larger domains, the unknowns \mathbf{s} are typically
58 discretized on a regular and equispaced grid. Under the common assumption that the
59 unknowns \mathbf{s} are statistically second order stationary (or intrinsic in such a way that an
60 adequate choice of \mathbf{X} yields $\mathbf{Q}_{\mathbf{ss}}$ second-order stationary), $\mathbf{Q}_{\mathbf{ss}}$ has symmetric Toeplitz
61
62
63
64
65

1
2
3 structure. In Toeplitz matrices, the entries along each diagonal have the same value.
4
5 Hence, the $(i + 1)$ -th row equals the i -th row translated by one entry to the right, with
6
7 a new element appearing on the first position on the left. In the d -dimensional case, the
8
9 \mathbf{Q}_{ss} has level- d block-Toeplitz structure, where the same pattern applies to nested block
10
11 structures.
12

13
14 As a consequence, the first column of a symmetric Toeplitz matrix contains all the in-
15
16 formation. This has been shown to be computationally exploitable (Zimmerman 1989).
17
18 First of all, only one column has to be stored instead of a full matrix, reducing storage
19
20 requirements from n^2 to n elements. Second, many algorithms have been found that work
21
22 on the first column only, or that work on the generating vectors only (e.g., Kailath and
23
24 Sayed 1995).
25

26
27 A Toeplitz matrix with a first column in the form of $a_0, a_1, \dots, a_N, \dots, a_1$ is called a circulant
28
29 matrix. Any Toeplitz matrix \mathbf{Q}_{ss} with first column q_0, q_1, \dots, q_n can be embedded in a larger
30
31 circulant matrix $\mathbf{Q}_{\text{ss},e}$. The physical analog is embedding a finite domain in a larger
32
33 periodic one. The easiest way to do this is to simply append the elements q_1, \dots, q_{n-1}
34
35 in reverse order $q_{n-1}, q_{n-2}, \dots, q_1$ to the first column. For specific applications, smaller
36
37 embedding sizes may be allowable, or larger ones to ensure positive-definiteness may be
38
39 required. For example, the generation of random fields require the resulting circulant
40
41 matrix to be non-negative, and it should be positive-definite if used together with certain
42
43 iterative solvers, while convolution via FFT unaffected by the definiteness. These issues
44
45 are discussed in more detail by Newsam and Dietrich (1994), Dietrich and Newsam (1997),
46
47 or Nowak et al. (2003).
48
49

50 51 Embedding and Extraction

52
53 In matrix notation, embedding and extraction may be formally expressed via an $n_e \times n$
54
55 mapping matrix \mathbf{M} (Cirpka and Nowak (2004)):
56
57

$$58 \quad \mathbf{M} = \begin{bmatrix} \mathbf{I}_{n \times n} \\ \mathbf{0}_{(n_e - n) \times n} \end{bmatrix} \quad (9)$$

1
2
3 Then, embedding an $n \times 1$ vector \mathbf{x} is
4
5
6

$$7 \quad \text{embedding: } \mathbf{x}^e = \mathbf{M}\mathbf{x} \quad (10)$$

8
9
10 and extracting from an $n_e \times 1$ vector \mathbf{x}_e is
11
12
13

$$14 \quad \text{extraction: } \mathbf{x} = \mathbf{M}^T \mathbf{x}_e \quad (11)$$

15
16
17
18 The Toeplitz matrix can formally be extracted from the embedding circulant one by
19
20
21

$$22 \quad \mathbf{Q}_{\text{ss}} = \mathbf{M}^T \mathbf{Q}_{\text{ss},e} \mathbf{M} . \quad (12)$$

23
24
25
26 Of course, embedding and extraction are achieved via zero-padding or disregarding ex-
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
cess elements instead of performing actual matrix-vector products, or by using a sparse
representation of \mathbf{M} . Throughout the remaining paper, \mathbf{Q}_{ss} denotes Toeplitz-structured
covariance matrices and $\mathbf{Q}_{\text{ss},e}$ denotes larger circulant covariance matrices with first col-
umn $\mathbf{q}_{\text{ss},e}$ that embed \mathbf{Q}_{ss} . Subscript e denotes vectors embedded according to Eq. (10).

66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127 128 129 130 131 132 133 134 135 136 137 138 139 140 141 142 143 144 145 146 147 148 149 150 151 152 153 154 155 156 157 158 159 160 161 162 163 164 165 166 167 168 169 170 171 172 173 174 175 176 177 178 179 180 181 182 183 184 185 186 187 188 189 190 191 192 193 194 195 196 197 198 199 200 201 202 203 204 205 206 207 208 209 210 211 212 213 214 215 216 217 218 219 220 221 222 223 224 225 226 227 228 229 230 231 232 233 234 235 236 237 238 239 240 241 242 243 244 245 246 247 248 249 250 251 252 253 254 255 256 257 258 259 260 261 262 263 264 265 266 267 268 269 270 271 272 273 274 275 276 277 278 279 280 281 282 283 284 285 286 287 288 289 290 291 292 293 294 295 296 297 298 299 300 301 302 303 304 305 306 307 308 309 310 311 312 313 314 315 316 317 318 319 320 321 322 323 324 325 326 327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343 344 345 346 347 348 349 350 351 352 353 354 355 356 357 358 359 360 361 362 363 364 365 366 367 368 369 370 371 372 373 374 375 376 377 378 379 380 381 382 383 384 385 386 387 388 389 390 391 392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408 409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425 426 427 428 429 430 431 432 433 434 435 436 437 438 439 440 441 442 443 444 445 446 447 448 449 450 451 452 453 454 455 456 457 458 459 460 461 462 463 464 465 466 467 468 469 470 471 472 473 474 475 476 477 478 479 480 481 482 483 484 485 486 487 488 489 490 491 492 493 494 495 496 497 498 499 500 501 502 503 504 505 506 507 508 509 510 511 512 513 514 515 516 517 518 519 520 521 522 523 524 525 526 527 528 529 530 531 532 533 534 535 536 537 538 539 540 541 542 543 544 545 546 547 548 549 550 551 552 553 554 555 556 557 558 559 560 561 562 563 564 565 566 567 568 569 570 571 572 573 574 575 576 577 578 579 580 581 582 583 584 585 586 587 588 589 590 591 592 593 594 595 596 597 598 599 600 601 602 603 604 605 606 607 608 609 610 611 612 613 614 615 616 617 618 619 620 621 622 623 624 625 626 627 628 629 630 631 632 633 634 635 636 637 638 639 640 641 642 643 644 645 646 647 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695 696 697 698 699 700 701 702 703 704 705 706 707 708 709 710 711 712 713 714 715 716 717 718 719 720 721 722 723 724 725 726 727 728 729 730 731 732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770 771 772 773 774 775 776 777 778 779 780 781 782 783 784 785 786 787 788 789 790 791 792 793 794 795 796 797 798 799 800 801 802 803 804 805 806 807 808 809 810 811 812 813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829 830 831 832 833 834 835 836 837 838 839 840 841 842 843 844 845 846 847 848 849 850 851 852 853 854 855 856 857 858 859 860 861 862 863 864 865 866 867 868 869 870 871 872 873 874 875 876 877 878 879 880 881 882 883 884 885 886 887 888 889 890 891 892 893 894 895 896 897 898 899 900 901 902 903 904 905 906 907 908 909 910 911 912 913 914 915 916 917 918 919 920 921 922 923 924 925 926 927 928 929 930 931 932 933 934 935 936 937 938 939 940 941 942 943 944 945 946 947 948 949 950 951 952 953 954 955 956 957 958 959 960 961 962 963 964 965 966 967 968 969 970 971 972 973 974 975 976 977 978 979 980 981 982 983 984 985 986 987 988 989 990 991 992 993 994 995 996 997 998 999 1000

Operations similar to embedding and extraction are injection and sampling. Measurement
as a process of sampling from a vector has been defined in Eq. 2, and injection is the
opposite operation. As formal matrix operations, sampling and injection are denoted by:

$$50 \quad \text{sampling: } \mathbf{x} = \mathbf{H}\mathbf{X} \quad (13)$$

$$51 \quad \text{injection: } \mathbf{X} = \mathbf{H}^T \mathbf{x} \quad (14)$$

52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000

1
2
3 methods to irregular grids.
4
5

6 Convolution via FFT 7

8
9 One of the most interesting properties of Toeplitz matrices is that multiplication with
10 a vector is the same as a discrete convolution of the vector with the first column of a
11 Toeplitz matrix. This simple and well-known trick is the basic building block for all
12
13
14
15
16 further efficient spectral methods discussed in this paper.
17

18 Consider \mathbf{x} an $n \times 1$ vector, \mathbf{Q}_{ss} a $n \times n$ Toeplitz matrix, $\mathbf{q}_{ss,e}$ the first column of the
19 embedding circulant matrix, and let $\mathcal{F}[\cdot]$ and $\mathcal{F}^{-1}[\cdot]$ denote the Fourier transform and
20 its inverse, respectively. Then:
21
22
23

$$24 \mathbf{Q}_{ss}\mathbf{x} = \mathbf{M}^T \mathcal{F}^{-1} [\mathcal{F} [\mathbf{M}\mathbf{x}] \circ \mathcal{F} [\mathbf{q}_{ss,e}]] , \quad (15)$$

25
26
27
28
29 where $(\cdot) \circ (\cdot)$ denotes the element-wise (Hadamard) product. Using the fast Fourier
30 transform (FFT) or its extension to arbitrary vector length (Cooley and Tukey 1965,
31 Frigo and Johnson 1998), Fourier transforming an $n \times 1$ vector has a computational com-
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
plexity of $\mathcal{O}(n \log_2 n)$. Taking a small detour into a larger periodic system via circulant
embedding, the product $\mathbf{Q}_{ss}\mathbf{x}$ can be evaluated efficiently while only storing the first
column of the matrix (e.g. van Loan 1992). The computational complexity is reduced
from $\mathcal{O}(n^2)$ to $\mathcal{O}(n \log n)$ and storage requirements are reduced from $\mathcal{O}(n^2)$ to $\mathcal{O}(n)$.
The storage requirements for standard evaluation of $\mathbf{Q}_{ss}\mathbf{x}$ may also be reduced to $\mathcal{O}(n)$
by n successive additions of the shifted integral kernel of convolution without storing
more than the current shifted version. The shifting procedure, however, also requires an
embedding procedure and the same number of memory access operations as the explicit
brute-force matrix product. The same FFT-based procedure and some extensions are
used for evaluating auto- and cross-covariance matrices of dependent quantities in the
context of cokriging and geostatistical inverse modeling in Nowak et al. (2003)).

Figure 1 (solid line and bold dash-dotted line) shows a comparison of CPU time and
storage requirements depending on vector length n on a contemporary desktop computer

1
2
3 (i386, 2.8GHz Intel Xeon dual-core, 2GB RAM, Suse Linux 9.2, implemented in MATLAB
4
5 R2006b). Although convolution via FFT has a certain overhead for embedding, it is faster
6
7 even for very small problems since only n instead of n^2 elements of \mathbf{Q}_{ss} need to be generated
8
9 and handled in memory. The costs of conventional convolution are included as bold dash-
10
11 dotted line for $n/m = 1$, while the other ratios of n over m refer to the superposition
12
13 problem in the following section. Convolution via FFT is faster by almost three orders of
14
15 magnitude unless for very small problems, where lower orders of computational complexity
16
17 play a role. The X-marked circles indicate memory overflow. For simple conventional
18
19 superposition, where n^2 elements need to be stored, storage is the main restriction to
20
21 the problem size. Successive shifting of the integral kernel remedies this restriction, but
22
23 does not save computational costs, resulting in the straight-line extension of the standard
24
25 methods beyond the point of memory overflow.
26
27

28 Superposition via FFT

29
30
31
32 Evaluating the term $\mathbf{Q}_{sy}\boldsymbol{\xi}$ Eq. (6) is identical to a $\boldsymbol{\xi}$ -weighted superposition of the auto-
33
34 covariance of the unknowns, where the individual offset between the single superimposed
35
36 terms is given by the measurement locations. In a formal step which at first looks counter-
37
38 productive, we show how to trace this operation back to a convolution using Eq. (4):
39
40
41

$$42 \quad \mathbf{Q}_{sy}\boldsymbol{\xi} = \mathbf{Q}_{ss}\mathbf{H}^T\boldsymbol{\xi} = \mathbf{Q}_{ss}\underbrace{(\mathbf{H}^T\boldsymbol{\xi})}_{\boldsymbol{\Xi}}. \quad (16)$$

43
44
45
46
47 The parentheses on the right-hand side of this equation points out that we first inject the
48
49 $m \times 1$ vector $\boldsymbol{\xi}$ into an $n \times 1$ vector $\boldsymbol{\Xi}$. The remaining matrix-vector product $\mathbf{Q}_{ss}\boldsymbol{\Xi}$ is
50
51 again a convolution, so that:
52
53

$$54 \quad \mathbf{Q}_{sy}\boldsymbol{\xi} = \mathbf{M}^T\mathcal{F}^{-1}[\mathcal{F}[\mathbf{M}\mathbf{H}^T\boldsymbol{\xi}] \circ \mathcal{F}[\mathbf{q}_{ss,e}]] , \quad (17)$$

55
56
57
58
59 in which we combine the injection and embedding of $\boldsymbol{\xi}$ into one operation. At a com-
60
61 putational complexity of $\mathcal{O}(n \log n)$, this scheme is faster than direct superposition for
62
63 $m > \log n$, i.e., for relatively large numbers of measurements. We also avoid the effort of
64
65

1
2
3 cutting/pasting /shifting the covariance function to different locations prior to summation
4
5 or, even worse, to store all of these shifted functions simultaneously. Instead, only the
6
7 first column $\mathbf{q}_{ss,e}$ needs to be stored, keeping the storage requirements at $\mathcal{O}(n)$ instead of
8
9 $\mathcal{O}(nm)$.

10
11 The resulting computational effort and storage requirements are shown in Figure 1 for
12
13 different values of n/m . In spite of its overhead and independence of m , the FFT-based
14
15 algorithm is faster than the standard method in the explored range of n/m . The largest
16
17 speedup, nearly three orders of magnitude compared to conventional superposition, is
18
19 achieved for high numbers m of injections relative to n .

22 23 **Efficient solution of Toeplitz systems**

24
25
26
27 Eqs. (7) and (8) require the solutions of systems using the $m \times m$ auto-covariance matrix
28
29 of the measurements, \mathbf{Q}_{yy} . If the measurements are arranged on a regular grid, \mathbf{Q}_{yy} has a
30
31 Toeplitz structure just like \mathbf{Q}_{ss} . In that case, we suggest solving the Kriging equations via
32
33 the FFT-based Preconditioned Conjugate Gradient (FFT-PCG) solver described below.
34
35 Other techniques, with slightly lower performance and less adequate for combination with
36
37 the superposition task, include look-ahead Schur algorithms and algorithms based on
38
39 generalized displacement structures (e.g. Gallivan et al. 1996, Kailath and Sayed 1999,
40
41 Barel et al. 2001) and are not discussed here.

42
43
44 The preconditioned conjugate gradient (PCG) method (Shewchuk 1994) iteratively solves
45
46 a linear system of equations $\mathbf{Ax} = \mathbf{b}$. The PCG method converges in as many iteration
47
48 steps as there are distinct eigenvalues of \mathbf{A} . If there exists a preconditioning matrix \mathbf{V}
49
50 which clusters the eigenvalues of the product $\mathbf{V}^{-1}\mathbf{A}$ around unity, the algorithm will
51
52 converge in only a few steps.

53
54
55 The advantage in this context is that PCG only requires evaluation of the product \mathbf{Ax} ,
56
57 which may be performed without explicitly storing or knowing all elements of \mathbf{A} . The
58
59 same holds for applying the preconditioner. In the FFT-PCG algorithm, these steps are
60
61 evaluated through convolution via FFT, i.e. in $\mathcal{O}(n \log_2 n)$ operations. Solving a circulant
62
63 system is identical to convolution via FFT with only the Hadamard product replaced by
64
65

1
2
3 element-wise division (compare Good 1950, Rino 1970). If choosing a circulant matrix as
4
5 preconditioner for the Toeplitz system, then applying the preconditioner is once again an
6
7 $n \log_2 n$ operation, and only one column of the preconditioner needs to be stored.
8
9

10 Chan and Ng (1996) review and compare circulant preconditioners that have the same
11
12 size as the original Toeplitz matrix and differ in the norm of $\mathbf{V}^{-1}\mathbf{A}$ they satisfy. For
13
14 poor-conditioned Toeplitz systems, e.g., for Gaussian covariance matrices with high cor-
15
16 relation lengths (e.g., Wesson and Pogram (2004)), the preconditioners themselves have
17
18 a very poor condition, causing numerical noise to be amplified. Wesson and Pogram
19
20 (2004) showed that singular value decomposition is very powerful in suppressing numeri-
21
22 cal artifacts when Kriging with such poorly conditioned systems. Trapp (1973) discussed
23
24 generalized inverses of circulant matrices which could be used in this context, but we
25
26 encountered poor convergence of generalized inverses with our iterative solver in prelim-
27
28 inary studies. Along similar lines, Nowak (2005) proposed a regularization term, which
29
30 is added to the diagonal of circulant preconditioners and evenly increases all eigenvalues
31
32 (Appendix) instead of disregarding the almost-zero ones. We chose the Strang precon-
33
34 ditioner (Strang 1986) on the embedded level, which is identical to the circulant matrix
35
36 already in use for convolution via FFT, in conjunction with the regularization suggested
37
38 by Nowak. However, if measurement error is included so that $\mathbf{R} \neq \mathbf{0}$, then \mathbf{R} regular-
39
40 izes the entire system anyway and it is unlikely that additional regularization is required
41
42 (compare Dietrich and Newsam 1989).
43
44
45

46 Figure 2 compares CPU times and storage requirements for standard Gaussian elimina-
47
48 tion (built-in solver for dense systems in MATLAB R2006b) versus the FFT-based PCG
49
50 algorithm for the solution of Toeplitz systems of different sizes m . For a number of un-
51
52 knowns of $m \approx 300$, a break-even point of the two algorithms can be observed. Still below
53
54 $m = 1,000$, FFT-based PCG is faster than the standard solver by more than an order
55
56 of magnitude. At $m = 10,000$, the standard solver broke down due to memory overflow,
57
58 whereas the FFT-based PCG worked up to $m = 1.6 \cdot 10^7$. The figure also demonstrates
59
60 the superiority of the FFT-PCG over other standard solvers equipped with FFT-based al-
61
62 gorithms for handling the Toeplitz matrix. Iterative Constrained Deconvolution (Pogram
63
64
65

2004, see also next section) is still faster than steepest descent (e.g., Press et al. 1992). The largest individual speedup is achieved by preconditioning, which requires more effort per iteration step but drastically reduces the number of iterations. The conjugation of gradients increases the effort per step once more, but the net effect of even fewer iteration steps prevails.

Extension of Toeplitz solvers to irregular grids

In order to solve Eq. 7 for irregularly scattered measurements, Pegram (2004) uses an algorithm called Iterative Constrained Deconvolution (ICD). The ICD algorithm can be identified as a steepest descent algorithm of the form $\xi_{i+1} = \xi_i + \alpha (\mathbf{Y} - \mathbf{Q}_{yy}\xi_i)$, where α denotes a user-defined step size coefficient. For the matrix-vector product $\mathbf{Q}_{yy}\xi_i$, he uses an injection/sampling combination

$$\mathbf{Q}_{yy}\xi_i = \mathbf{H} [\mathbf{Q}_{ss} (\mathbf{H}^T \xi_i)] \quad (18)$$

and evaluates the product with \mathbf{Q}_{ss} using convolution via FFT. To be more precise, the sampling (multiplication with \mathbf{H}) is never actually performed during the algorithm. Instead, the final value within the square brackets is used for $\mathbf{Q}_{sy}\xi$ in Eq. 6, thus avoiding one additional superposition.

In order to extend the FFT-based PCG to irregular sampling grids, we seized Pegram's suggestion to inject the irregularly scattered measurements into a finer regular grid, as also done in Eq. 18. The corresponding injection/sampling procedure $\mathbf{Q}_{yy} = \mathbf{H}\mathbf{Q}_{ss}\mathbf{H}^T$, where \mathbf{Q}_{ss} is a larger Toeplitz matrix, is formally identical to the embedding/extraction procedure $\mathbf{Q}_{ss} = \mathbf{M}^T\mathbf{Q}_{ss,e}\mathbf{M}$ to embed a Toeplitz matrix in a larger circulant matrix, and both procedures can be combined into one:

$$\begin{aligned} \mathbf{Q}_{yy} &= \mathbf{H}\mathbf{Q}_{ss}\mathbf{H}^T \\ \mathbf{Q}_{ss} &= \mathbf{M}^T\mathbf{Q}_{ss,e}\mathbf{M} \\ \rightarrow \mathbf{Q}_{yy} &= (\mathbf{H}\mathbf{M}^T) \mathbf{Q}_{ss,e} (\mathbf{M}\mathbf{H}^T) \end{aligned} \quad (19)$$

1
2
3 All other particulars are identical to the FFT-based PCG algorithm for Toeplitz matrices
4 (regular measurement grids). Placed within the PCG framework, the idea of sam-
5 pling/injection is much more efficient than within the framework of steepest descent as
6 used in the ICD algorithm. The resulting FFT-based PCG solver for almost-Toeplitz
7 matrices (irregularly scattered measurements) is described in Appendix.
8
9

10
11
12
13
14 Figure 3 compares CPU times and storage requirements for FFT-based solvers for nearly-
15 Toeplitz systems with those of standard Gaussian elimination for different problem size
16 m and sizes of the finer regular grid n . The finer grid introduces a substantial overhead if
17 a high spatial accuracy of measurement locations on the finer grid is desired. Therefore,
18 the standard solver is faster than the FFT-PCG in the majority of cases. However, the
19 reduced storage requirements allow FFT-based algorithms to solve much larger systems
20 of equations. On our reference computer, the limiting size of the underlying regular grid
21 was $n = 1.6 \cdot 10^7$. The FFT-PCG solver proposed in this study outperforms the ICD
22 algorithm by a factor of roughly ten.
23
24
25
26
27
28
29
30
31
32

33 34 APPLICATION OF FFT-BASED ALGORITHMS 35 36 37 TO KRIGING 38

39 40 41 Kriging Estimator 42 43 44

45 In order to apply the efficient solvers discussed above for solving the Kriging system of
46 equations (Eq. 7), we exploit the fact that the coefficient matrix (also called the Kriging
47 matrix) consists of an $m \times m$ sub-matrix \mathbf{Q}_{yy} with exploitable structure and structural
48 perturbations of rank $p \ll m$. First, we introduce three auxiliary variables which will be
49 re-used several times in the following procedures:
50
51
52
53
54
55
56

$$57 \quad \mathbf{x} = \mathbf{H}\mathbf{X} \quad (20)$$

$$58 \quad \mathbf{y} = \mathbf{Q}_{yy}^{-1}\mathbf{Y} \quad (21)$$

$$59 \quad \mathbf{z} = \mathbf{Q}_{yy}^{-1}\mathbf{H}\mathbf{X} = \mathbf{Q}_{yy}^{-1}\mathbf{x} \quad (22)$$

Then, we partition the inverse of the Kriging matrix in 7 as follows (Kitanidis 1996, Nowak and Cirpka 2004):

$$\begin{bmatrix} \mathbf{Q}_{yy} & \mathbf{HX} \\ \mathbf{X}^T \mathbf{H}^T & -\mathbf{Q}_{\beta\beta}^{-1} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{P}_{yy} & \mathbf{P}_{y\beta} \\ \mathbf{P}_{\beta y} & \mathbf{P}_{\beta\beta} \end{bmatrix} \quad (23)$$

Now, we express the sub-matrices \mathbf{P} in Eq. (23) according to Schwappe (1973) and immediately simplify using our auxiliary quantities:

$$\mathbf{P}_{\beta\beta} = -(\mathbf{x}^T \mathbf{z} + \mathbf{Q}_{\beta\beta}^{-1})^{-1} \quad (24)$$

$$\mathbf{P}_{\beta y} = \mathbf{P}_{y\beta}^T = -\mathbf{P}_{\beta\beta} \mathbf{z}^T \quad (25)$$

$$\mathbf{P}_{yy} = \mathbf{Q}_{yy}^{-1} + \mathbf{z} \mathbf{P}_{\beta\beta} \mathbf{z}^T \quad (26)$$

where Eq. 24 is also known as the Schur complement. Using these sub-matrices yields a partitioned form of the coefficient vector (compare Nowak and Cirpka 2004):

$$\begin{aligned} \boldsymbol{\xi} &= \mathbf{y} - \mathbf{z} \hat{\boldsymbol{\beta}} \\ \hat{\boldsymbol{\beta}} &= -\mathbf{P}_{\beta\beta} (\mathbf{z}^T \mathbf{Y} + \mathbf{Q}_{\beta\beta}^{-1} \boldsymbol{\beta}^*) \end{aligned} \quad (27)$$

Altogether, the entire estimator requires to:

1. Compute the auxiliary quantities \mathbf{x} , \mathbf{y} and \mathbf{z} according to Eqs. (20) to (22), using the most appropriate solver (e.g. the FFT-based PCG),
2. Evaluate the partitioned solution vector of the Kriging system according to Eq. (27),
3. Evaluate the estimate according to Eq. (6), using superposition via FFT (Eq. 17) for $\mathbf{Q}_{sy} \boldsymbol{\xi}$ and simple matrix-vector multiplication for $\mathbf{X} \hat{\boldsymbol{\beta}}$

Step one requires p sampling processes to obtain \mathbf{x} , and $1+p$ solutions of a (nearly) Toeplitz system to obtain \mathbf{y} and \mathbf{z} , with an asymptotic cost estimate of $\mathcal{O}((1+p)m \log_2 m)$ (when \mathbf{Q}_{ss} is second-order stationary and the measurements \mathbf{y} fall onto a regular grid so that the FFT-PCG for regular grids can be applied). Step two involves only smaller operations of

1
2
3 $\mathcal{O}(mp)$ and $\mathcal{O}(p^2)$ to treat the rank p perturbations in the structure of the Kriging matrix.
4
5 Step three is an $\mathcal{O}(np + n \log_2 n)$ operation. This yields a total of $\mathcal{O}(n \log_2 n + m \log_2 m)$
6
7 for the entire estimation with storage requirements of only $\mathcal{O}(n)$, given that $n \gg m \gg p$.
8
9

10 Estimation Variance

11
12 To obtain an efficient procedure for evaluating the estimation variance, we re-combine
13
14 some terms in Eq. (8):
15
16
17
18

$$19 \mathbf{Q}_{ss|y} = \mathbf{Q}_{ss} - \underbrace{\begin{bmatrix} \mathbf{Q}_{ys} \\ \mathbf{X}^T \end{bmatrix}^T \begin{bmatrix} \mathbf{Q}_{yy} & \mathbf{HX} \\ (\mathbf{HX})^T & -\mathbf{Q}_{\beta\beta}^{-1} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{Q}_{ys} \\ \mathbf{X}^T \end{bmatrix}}_{\mathbf{S}} \quad (28)$$

$$20 \mathbf{S} = \begin{bmatrix} \mathbf{Q}_{ys} \\ \mathbf{X}^T \end{bmatrix}^T \begin{bmatrix} \mathbf{Q}_{yy} & \mathbf{HX} \\ (\mathbf{HX})^T & -\mathbf{Q}_{\beta\beta}^{-1} \end{bmatrix}^{-1} \mathbf{I}_{(m+p)} \quad (29)$$

21
22 Here, we appended the $(m+p) \times (m+p)$ identity matrix \mathbf{I} to \mathbf{S} without actually changing
23
24 the equation. However, it clarifies that the $(m+p)$ columns of \mathbf{S} are unit estimators \mathbf{s}_i
25
26 that arise from using the $(m+p)$ unit vectors \mathbf{e}_i in $\mathbf{I}_{(m+p)}$ as data vectors for Kriging in
27
28 Eq. (7).
29
30
31
32

33 For the estimation variance $\hat{\sigma}^2$, we only need the diagonal of $\mathbf{Q}_{ss|y}$. The diagonal of a
34
35 matrix product or of dyadic matrices (i.e., matrices defined by dyadic products) with low
36
37 rank can be directly and efficiently evaluated using
38
39
40
41

$$42 \text{diag}(\mathbf{A}_{(n \times m)} \mathbf{B}_{(n \times m)}^T) = \sum_{i=1}^m (\mathbf{a}_i \circ \mathbf{b}_i) \quad (30)$$

43 where \mathbf{a}_i and \mathbf{b}_i are the i -th columns of \mathbf{A} and \mathbf{B} , respectively. Then, Eq. (28) simplifies
44
45 to:
46
47
48
49

$$50 \hat{\sigma}^2 = \text{diag}(\mathbf{Q}_{ss|y}) = \sigma^2 - \sum_{i=1}^{m+p} \mathbf{s}_i \circ \underbrace{\begin{bmatrix} \mathbf{Q}_{sy} & \mathbf{X} \end{bmatrix}_i}_{\sigma_{kriging}^2} \quad (31)$$

To efficiently evaluate this expression:

1. Initialize $\hat{\sigma}^2 = \sigma^2$,
2. Evaluate the unit estimator \mathbf{s}_i for the $(m + p) \times 1$ unit vector \mathbf{e}_i as data vector for Kriging, using FFT-based algorithms and Eqs. (20) to (27) as described in the previous section,
3. Perform the Hadamard product of the unit estimator with the i -th column of $\begin{bmatrix} \mathbf{Q}_{\text{sy}} & \mathbf{X} \end{bmatrix}$ and subtract the result from $\hat{\sigma}^2$ and
4. Repeat steps 2 and 3 for $i = 1 \dots m + p$.

In total, this makes evaluating $\hat{\sigma}^2$ an effort of $(m + p)$ estimations, $(m + p)$ Hadamard products and summation of $(m + p)$ vectors, each sized $n \times 1$, resulting in an asymptotic complexity of $\mathcal{O}(mn \log_2 n + m^2 \log_2 m)$.

Efficient Approximations to the Estimation Variance

For the sake of the following analysis, we revert to the partitioned form of the Kriging matrix. Starting from Eq. (8), inserting Eqs. 24 to 26 and applying Eq. 30 yields after some rearrangement:

$$\begin{aligned} \sigma_{kriging}^2 = & \sum_{i=1}^m [\mathbf{Q}_{\text{sy}} \mathbf{Q}_{\text{yy}}^{-1}]_i \circ \mathbf{Q}_{\text{sy},i} + \sum_{i=1}^p [\mathbf{Q}_{\text{sy}} \mathbf{zP}_{\beta\beta}]_i \circ [\mathbf{Q}_{\text{sy}} \mathbf{z}]_i \\ & - 2 \sum_{i=1}^p [\mathbf{Q}_{\text{sy}} \mathbf{zP}_{\beta\beta}]_i \circ \mathbf{X}_i + \sum_{i=1}^p [\mathbf{XP}_{\beta\beta}]_i \circ \mathbf{X}_i \end{aligned} \quad (32)$$

The first term is the one required for Kriging with known mean, i.e. with a number of base functions $p = 0$, and represents m unit estimators similar to those defined in Eq. 29. Hence, it has well-known properties. The approximations in the following sections will focus on the asymptotic behavior of this first term under specific conditions. With no further simplifications, this term requires full inversion of \mathbf{Q}_{yy} , m superpositions via FFT to evaluate $\mathbf{Q}_{\text{sy}} \mathbf{Q}_{\text{yy}}^{-1}$ and then m Hadamard vector-vector products, followed by summation over index i .

All other terms relate to the uncertainty in estimating the trend coefficients. We will not simplify them any further, since their computational complexity is negligible over that of the first term, given that $m \gg p$: We first evaluate p superpositions of $\mathcal{O}(n \log_2 n)$ each to obtain the individual columns of the $n \times p$ matrix $\zeta = \mathbf{Q}_{\text{sy}} \mathbf{z}$ via FFT. The remaining steps are $\mathcal{O}(np^2)$ for the multiplication with $\mathbf{P}_{\beta\beta}$ and the $3p$ Hadamard products.

Single-Point Approximations

In the single-point approximations discussed here, correlation between measurements is neglected. They have been suggested by Pegram (2004) for zero measurement error and the specific cases of known/unknown mean. These approximations reduce the computational complexity of the estimation variance by one order in m .

1. For vanishing correlation among the measurement points on both irregular and regular grids, \mathbf{Q}_{yy} approaches a scalar matrix $\mathbf{Q}_{\text{yy}} \approx (\sigma^2 + \sigma_{\text{err}}^2) \mathbf{I}$ with σ_{err}^2 being the variance of measurement error. Then, the first term can be approximated by:

$$\sum_{i=1}^m [\mathbf{Q}_{\text{sy}} \mathbf{Q}_{\text{yy}}^{-1}]_i \circ \mathbf{Q}_{\text{sy},i} \approx \frac{1}{\sigma^2 + \sigma_{\text{err}}^2} \sum_{i=1}^m \mathbf{Q}_{\text{sy},i} \circ \mathbf{Q}_{\text{sy},i} \quad (33)$$

2. For the case of known-mean Kriging, the estimation variance is equal to the value of the measurement error variance σ_{err}^2 at measurement locations. If the correlation among measurements is too large for the above simplification, the latter condition may still be enforced by requiring that the first term equals $\sigma^2 - \sigma_{\text{err}}^2$ at the locations of measurements. This is achieved by solving the subsidiary Kriging problem:

$$\sum_{i=1}^m [\mathbf{Q}_{\text{sy}} \mathbf{Q}_{\text{yy}}^{-1}]_i \circ \mathbf{Q}_{\text{sy},i} \approx \mathbf{Q}_{\text{sy}}^* (\mathbf{Q}_{\text{yy}}^*)^{-1} \mathbf{u}_m (\sigma^2 - \sigma_{\text{err}}^2) \quad (34)$$

where \mathbf{u}_m is a $m \times 1$ vector of ones used as data vector in the subsidiary problem,

$$\mathbf{Q}_{\text{ss}}^* = \sigma^{-2} \mathbf{Q}_{\text{ss}} \circ \mathbf{Q}_{\text{ss}}, \quad \mathbf{Q}_{\text{sy}}^* = \mathbf{Q}_{\text{ss}}^* \mathbf{H}^T, \quad \text{and} \quad \mathbf{Q}_{\text{yy}}^* = \mathbf{H} \mathbf{Q}_{\text{ss}}^* \mathbf{H}^T.$$

Infinite Regular Grid Approximations

If the grid of measurements is regular and equispaced and its extent is much larger than the

1
2
3 range of correlation, the number of mutually correlated measurements differs only within
4
5 a thin boundary area. Besides that area, the inner section of the measurement behaves
6
7 statistically stationary. If the application permits, one may neglect these boundary effects
8
9 and assume that all measurement points have the same correlation to the same number of
10
11 neighbors. Then, the contribution of one single representative measurement point in the
12
13 center of the domain towards the estimation variance may be used to approximate that
14
15 of all others by shifting:
16
17
18

$$19 \quad \mathbf{Q}_{\text{sy}} \mathbf{Q}_{\text{yy}}^{-1} \approx \text{shift}_m [\mathbf{Q}_{\text{sy}} \boldsymbol{\xi}_c] \quad (35)$$

20
21
22
23 The term within square brackets is a $n \times 1$ zero-mean Representative Unit Estimator
24
25 (RUE) for a unit measurement in the center of the domain, and $\text{shift}_m [\cdot]$ denotes an
26
27 $n \times m$ matrix obtained from concatenating the RUE m times, each time shifted to the
28
29 respective next measurement location. To allow this shifting operation, the RUE needs
30
31 to be larger than the actual domain. This is achieved by performing the superposition
32
33 $\mathbf{Q}_{\text{sy}} \boldsymbol{\xi}_c$ for a virtual larger domain, using the embedded version of the covariance function.
34
35
36 Since measurements in the interior convey less independent information than those at
37
38 the boundaries, this approximation yields a conservative upper bound of the estimation
39
40 variance. It is exact for the case of an infinite (or periodic) grid of measurements. The
41
42 remaining complexity is again reduced by one order in m .
43
44

45 *Hybrid Regular Grid Approximation*

46
47
48 In the infinite regular grid approximation, the largest errors occur along the boundaries
49
50 of the measurement grid. The corresponding unit estimators may be evaluated separately
51
52 in their exact form, whereas the bulk inner part of the unit estimators are approximated
53
54 by shifting the RUE. A user-defined break criterion for the observed difference between
55
56 RUE and exact unit estimator can be used to determine the boundary zone in which to
57
58 use the exact version.
59
60
61
62
63
64
65

Another option to simplify the estimation variance is to start with a conditional covariance made stationary by averaging over all starting points of separation vectors with same length, i.e. by averaging each diagonal of $\mathbf{Q}_{\text{ss|y}}$. A semi-analytical solution for the averaged conditional covariance has been developed by Cirpka and Nowak (2003). We will translate the formalism into the notation and context used here for the sake of completeness. Due to averaging, the diagonals $\text{diag}(\cdot)$ in Eqs. (30) and (31) become traces $\text{Tr}(\cdot)$ multiplied by a factor of $1/n$. Since the trace is invariant with respect to cyclic permutations, the averaged equivalent of Eq. (32) is:

$$\begin{aligned} \overline{\sigma_{kriging}^2} = & \frac{1}{n} \text{Tr}(\mathbf{Q}_{\text{yy}}^{-1} \mathbf{Q}_{\text{ys}} \mathbf{Q}_{\text{sy}}) + \frac{1}{n} \text{Tr}(\mathbf{P}_{\beta\beta} \mathbf{z}^T \mathbf{Q}_{\text{ys}} \mathbf{Q}_{\text{sy}} \mathbf{z}) \\ & - \frac{2}{n} \text{Tr}(\mathbf{P}_{\beta\beta} \mathbf{X}^T \mathbf{Q}_{\text{sy}} \mathbf{z}) + \frac{1}{n} \text{Tr}(\mathbf{P}_{\beta\beta} \mathbf{X}^T \mathbf{X}) \end{aligned} \quad (36)$$

In this form, all terms but the first have boiled down to traces of $p \times p$ matrices. If we again evaluate $\zeta = \mathbf{Q}_{\text{sy}} \mathbf{z}$ via FFT, these small matrices are obtained from $3p$ scalar products and are computationally cheap.

For regular measurement grids, the first term is the inverse of an $m \times m$ Toeplitz matrix multiplied with another $m \times m$ Toeplitz matrix $\mathbf{Q}_{\text{ys}} \mathbf{Q}_{\text{sy}} = \mathbf{H} \mathbf{Q}_{\text{ss}} \mathbf{Q}_{\text{ss}} \mathbf{H}^T$. For small numbers of measurements, it may be evaluated exactly in a manner similar to the unit estimator technique described above, but the procedure may quickly explode in computational effort for larger m . If the measurement grid is sufficiently large or periodic, we may approximate both matrices by circulant counterparts, and evaluate first column of the circulant resulting from the matrix product by convolution via FFT of the individual first columns. The trace of the resulting circulant is quite trivially given by m times the first element (e.g., Davis 1979).

Generation of unconditional and conditional random fields

The fast and exact generation of random fields via FFT has been described by Dykaar

1
2
3 and Kitanidis (1992) and by Dietrich and Newsam (1993). The generation of statistically
4 stationary random fields on regular grids is based on the symmetric decomposition of \mathbf{Q}_{ss}
5 achieved by taking the element-wise square root of its eigenvalues obtained from $\mathcal{F}[\mathbf{q}_{ss,e}]$,
6 followed by convolution via FFT with a white-noise vector. The overall computational
7 effort is as low as $\mathcal{O}(n \log_2 n)$ at storage requirements of, once again, only $\mathcal{O}(n)$. Con-
8 ditioning to direct or linearly dependent measurements has been dealt with by Dietrich
9 and Newsam (1996). For the non-linear case, the same procedure is suggested by Ki-
10 tanidis (1995). They correct unconditional realizations to conditional ones by subtracting
11 a Kriging estimator with adequately chosen measurements. Without going into further
12 details, we point out that the FFT-based algorithms for Kriging described in this study
13 can also be applied in this correction step, keeping the asymptotic cost estimate down at
14 $\mathcal{O}(mn \log_2 n + m^2 \log_2 m)$ and memory consumption at $\mathcal{O}(n)$.

30 PERFORMANCE TESTS

31
32
33 In this section, the performance of the described FFT-based methods is demonstrated
34 in comparison with standard methods. Our performance analysis was carried out on a
35 contemporary desktop computer (i386, 2.8GHz Intel Xeon dual-core, 2GB RAM, Suse
36 Linux 9.2) purchased in 2004. All methods were implemented in MATLAB (Release
37 2006b). As conventional solver for dense systems, the built-in Gaussian elimination in
38 MATLAB was used. MATLAB also includes the Basic Linear Algebra Package and an
39 implementation of the FFTW (Frigo and Johnson 1998). In order to keep the FFTW
40 algorithm efficient for arbitrary domain sizes, we implemented a small algorithm that
41 chooses embedding sizes with prime factors of 2, 3, 5 and 7 only. Required relative error
42 norm for all iterative solvers was set to 10^{-10} .

43
44
45 In all performance tests, we used random measurement data and varied both the number n
46 of estimation points and the number m of measurements. All tests assumed an uncertain
47 but spatially constant mean value, so that the number p of trend coefficients is one and \mathbf{X}
48 is a $n \times 1$ vector of ones. The individual test series are composed of $n = 2^k$, $k = 2 \dots 24$
49 estimation points and ratios between n and the number m of estimation points given by
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3 $n/m = 2^\ell$, $\ell = 2 \dots 14$. All computations above 10^5s (approx. 1 day) were stopped and
4
5 their CPU times estimated by extrapolation of fitted complexity models.
6
7

8 **Kriging with conventional solver**

9

10
11 Standard Kriging on a regular estimation grid and with irregularly scattered measure-
12
13 ments serves as our base-case for comparison. With standard Kriging, we refer to Gaus-
14
15 sian elimination for solving the Kriging system, conventional superposition by successive
16
17 addition of the shifted covariance kernel, and evaluation of the exact estimation variance
18
19 The results are displayed in Figure 4 (solid lines). The auto-covariance matrix of the
20
21 unknowns \mathbf{s} is represented by its first column only in our conventional superposition algo-
22
23 rithm. Hence, mainly the size of the Kriging matrix, i.e. the number m of measurements
24
25 squared, is the limiting factor to the problem size. Had we chosen a brute-force matrix
26
27 product approach to superposition and had we explicitly stored the $n \times m$ matrix \mathbf{Q}_{sy} ,
28
29 the limitation would have been even more severe. Due to the restriction in size, CPU
30
31 time never rose above one day. When using the one-point approximation to simplify the
32
33 estimation variance (not shown here), the computational effort reduces by a factor of
34
35 roughly m .
36
37
38
39

40 The speedup of superposition via FFT versus conventional superposition can be read
41
42 from the dashed lines in Figure 4. Since the overall complexity is dominated by the
43
44 solution of the system, the order of the asymptotic cost estimate does not change for high
45
46 measurement numbers, but speedup factors of up to fifty occur. The memory-related
47
48 restrictions of measurement numbers still holds, and the storage requirements of both
49
50 versions coincide.
51
52

53 **Kriging with FFT-based solvers**

54
55
56

57 Now, we demonstrate the effect of using the FFT-based PCG solver instead of Gaussian
58
59 elimination, combined with superposition via FFT (Figure 5). The FFT-based PCG solver
60
61 is slightly slower for unfavorable ratios n/m (compare Figure 3), but radically cancels the
62
63 dependence of storage requirements on the number of measurements. As a consequence,
64
65

1
2
3 only the grid of estimation limits the admissible problem size. For small numbers of
4
5 measurements, it is more efficient to revert to conventional Gaussian elimination (dashed
6
7 lines in Figure 4).

8
9
10 The CPU time for evaluating the estimation variance is $\mathcal{O}(m)$ larger than that of one
11
12 Kriging estimate. Figure 5 also shows how approximations to the estimation variance
13
14 can drastically reduce the overall computational effort by approximately this factor of m .
15
16 For large numbers of measurements (small n/m), the observed advantage in CPU time
17
18 rises up to five orders of magnitude. When evaluating only the estimator and omitting
19
20 the estimation variance (not shown here), CPU times would show the same shape of
21
22 dependence on n and n/m as the dashed lines, reduced by a factor of approximately two.
23
24

25 The greatest advantage can be made if the measurements are on a regular grid. For
26
27 that case, the FFT-based PCG solver outruns Gaussian elimination for any problem size
28
29 and permits vast numbers of measurements at even lower computational expense. The
30
31 CPU times are smaller by about one order of magnitude over a large range of problem
32
33 sizes (Figure 6). Especially for large problems, the infinite-grid approximation of the
34
35 estimation variance is likely to be sufficiently accurate for many purposes. The associated
36
37 computational speedup is again about ten compared to the irregular-grid case, and again
38
39 about 5 orders of magnitude compared to the exact evaluation of the estimation variance.
40
41 Entirely omitting the estimation variance (not shown here), as in the previous case, results
42
43 in a reduction of CPU times down to roughly one half of the dashed lines.
44
45
46

47 **SUMMARY AND CONCLUSIONS**

48
49
50
51 In this study we compiled a toolbox of existing and newly extended FFT-based methods
52
53 to drastically speed up Kriging for large estimation problems. These methods include
54
55 FFT-based convolution, FFT-based superposition and an FFT-based PCG solver. All
56
57 methods apply to estimation on regular and equispaced grids. The measurements may
58
59 either lie on regular and equispaced grids or be irregularly scattered. For irregularly
60
61 scattered measurements, we introduced an extension of the existing FFT-based PCG
62
63 method which can be understood as an extension of Pegram's (2004) method, where
64
65

1
2
3 scattered measurements are assigned to nodes of a (finer) regular grid. The density of
4
5 the regular grid is dictated by the required accuracy in discretizing the irregular grid,
6
7 depending on the specific application. We showed how to apply these methods to the
8
9 generalized case of Kriging with uncertain mean and trends, including measurement error.
10
11 The efficiency of all methods was demonstrated in a series of performance tests on an
12
13 ordinary desktop computer with all algorithms implemented in MATLAB.
14

15
16 The first eye-catching advantage of all FFT-based methods is that only the covariance
17
18 function for one point, i.e. the first column of the respective auto-covariance matrices,
19
20 needs to be stored. This makes the storage requirements shrink from $\mathcal{O}(n^2)$ and $\mathcal{O}(m^2)$
21
22 for whole covariance matrices to only $\mathcal{O}(n)$ and $\mathcal{O}(m)$ for one column each (where n
23
24 is the number of estimation points and m is the number of measurements). This also
25
26 reduces the time needed to set up the matrices themselves, which is not negligible since
27
28 evaluating effective separation distances and covariance functions (including exponentials)
29
30 has substantial computational costs.
31

32
33 FFT-based superposition has computational costs in the order of $n \log_2 n$ instead of mn ,
34
35 which turned out to be highly favorable even for very small data sets on our reference
36
37 computer. For regular and equispaced measurement grids, the FFT-based PCG solver
38
39 with its complexity of $\mathcal{O}(m \log_2 m)$ outruns standard solvers for dense systems by far. For
40
41 irregularly scattered measurements, the extended FFT-based PCG solver is less efficient
42
43 than the standard solver for low and medium numbers of measurements. The break-
44
45 even point is at roughly $m = 1000$ measurements, depending on the resolution of the
46
47 finer regular grid. The great advantage, however, is that the FFT-based solver does not
48
49 require to store the $m \times m$ coefficient matrix, so it is applicable for arbitrarily high m .
50
51 Any solver that requires to explicitly store the dense $m \times m$ Kriging broke down due to
52
53 memory restrictions at a maximum of $m = 10,000$ on our reference computer. Combining
54
55 these methods, the Kriging estimator can be evaluated up to several millions of estimation
56
57 points in no more than seconds up to a few minutes on contemporary desktop computers.
58
59
60
61 Evaluation of the estimation variance is computationally much more demanding. For the
62
63 evaluation of the exact estimation variance, an equivalent of m estimation procedures
64
65

1
2
3 has to be performed, which may turn out to be strictly inhibiting for larger data sets of
4
5 measurements, in spite of the speedup achieved by FFT-based methods. We alleviated
6
7 this situation by reviewing, extending and proposing several fast approximations which
8
9 are asymptotically exact in specific cases. These cases include negligible correlation among
10
11 the measurements and very large regular grids of measurements. The approximations offer
12
13 an additional speedup of up to five orders of magnitude for large m , so that the largest
14
15 admissible problem could be completed in less than one day.
16

17
18 In summary, the methods reviewed, discussed and introduced in this paper greatly re-
19
20 duce the computational effort and storage requirements of Kriging, allowing to handle
21
22 several millions of estimation points and thousands of measurements on ordinary desktop
23
24 computers.
25

26 27 28 **OUTLOOK** 29

30
31
32 For the near future, we expect further advances in the development of non-uniform FFT
33
34 algorithms (NUFFT), also called the generalized FFT (e.g., Liu and Ngyuen 1998, Dui-
35
36 jndam and Schonewille 1999, Fourmont 2003, Fessler and Sutton 2003, Greengard and
37
38 Lee 2004). The Kriging context requires NUFFT algorithms for frequency-space data
39
40 on a regular grid with real-space data on an irregular grid, that work efficiently in both
41
42 transform directions. We assume that highly efficient algorithms for these requirements
43
44 will be readily available within a few years. Such algorithms will help to further increase
45
46 the efficiency of preconditioning the PCG solver for irregularly scattered measurements.
47
48

49 50 51 **ACKNOWLEDGMENTS** 52

53
54 We would like to thank Prof. Pegram at the University of KwaZulu-Natal, South Africa,
55
56 for his helpful comments on the manuscript. This study was fundet in Parts by Deutsche
57
58 Forschungsgemeinschaft (DFG) under the grants Ne 824/2-2 and Ne 824/3.
59
60
61
62
63
64
65

APPENDIX

FFT-based PCG for irregular and regular grids

The Conjugate Gradients Method is attributed to Hestenes and Stiefel (1952). The following is the preconditioned version taken from Shewchuk (1994), combined with convolution via FFT (e.g. van Loan 1992) and the embedded regularized preconditioner by Nowak (2005). For an extensive review of circulant preconditioners used in PCG algorithms, see Chan and Ng (1996).

The Fourier transform of the first column of a circulant matrix yields the eigenvalues, thanks to the diagonalization theorem (e.g., Trapp 1973, Barnett 1990 pp. 350-354). If \mathbf{C} is a (level- d) real circulant matrix with first column \mathbf{c} and $\tilde{\mathbf{c}} = \mathcal{F}_d(\mathbf{c})$ is the (d -dimensional) Fourier transform of the first column, then the condition c of \mathbf{C} is the ratio of the largest value c_{max} and the smallest value c_{min} in $\tilde{\mathbf{c}}$. The regularized preconditioner by Nowak (2005) installs a maximum condition c^* (e.g. 10^5) of the preconditioner through a diagonal regularization $\mathbf{C}^* = \mathbf{C} + \varepsilon^* \mathbf{I}$, where ε is chosen according to

$$\varepsilon^* = \frac{c_{max} - c_{min} c^*}{c^* - 1}. \quad (37)$$

Automatically, the Fourier transform $\tilde{\mathbf{c}}^* = \mathcal{F}_d(\mathbf{c}^*)$ of the first column of \mathbf{C}^* is given by $\tilde{\mathbf{c}}^* = \tilde{\mathbf{c}} + \varepsilon^*$.

In the following, \circ denotes the element-wise (Hadamard) product and \div denotes element-wise division.

Algorithm 1 (*Preconditioned Conjugate Gradients with circulant preconditioning for nearly-Toeplitz system*): The linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ is to be solved for a real symmetric positive-definite $m \times m$ matrix $\mathbf{A} = \mathbf{H}\mathbf{T}\mathbf{H}^T$, where \mathbf{H} is a sampling matrix as defined in Eq. 2 and \mathbf{T} is a (level- d) symmetric positive-definite Toeplitz matrix sized $n \times n$. \mathbf{T} has a symmetric positive-definite embedding circulant matrix \mathbf{C} with $\mathbf{T} = \mathbf{M}^T \mathbf{C} \mathbf{M}$, where \mathbf{M} is a mapping matrix as defined in Eq. 9. \mathbf{c} is the first column of \mathbf{C} with (d -dimensional) Fourier transform $\tilde{\mathbf{c}} = \mathcal{F}_d(\mathbf{c})$. An initial guess

1
2
3 \mathbf{x}_0 , an error tolerance $\varepsilon < 1$ and a maximum allowable condition c^* are provided.
4
5 Initialize the algorithm with counter $k = 1$, error vector $\mathbf{r} = \mathbf{b} - \mathbf{A}\mathbf{x}_0$, the precon-
6
7 ditioned conjugate gradient $\mathbf{d} \approx \mathbf{A}^{-1}\mathbf{r}$, the residual $\delta_1 = \mathbf{r}^T\mathbf{d}$, the initial residual
8
9 $\delta_0 = \delta_1$ and evaluate ε^* according to Eq. (37). Then,
10

11
12 1. update the trial solution \mathbf{x} using:
13

$$\begin{aligned} 14 \quad \mathbf{q} &= \mathbf{H}\mathbf{M}^T \mathcal{F}_d^{-1} [\mathcal{F}_d [\mathbf{M}\mathbf{H}^T \mathbf{d}] \circ \tilde{\mathbf{c}}] \quad (= \mathbf{A}\mathbf{d}) \\ 15 \\ 16 \quad \alpha &= \frac{\delta_k}{\mathbf{d}^T \mathbf{q}} \\ 17 \\ 18 \quad \mathbf{x} &= \mathbf{x} + \alpha \mathbf{d} \\ 19 \\ 20 \\ 21 \\ 22 \\ 23 \\ 24 \end{aligned}$$

25 2. update the error vector and residual:
26

$$\begin{aligned} 27 \\ 28 \\ 29 \quad \mathbf{r} &= \mathbf{r} - \alpha \mathbf{q} \\ 30 \\ 31 \quad \mathbf{s} &= \mathbf{H}\mathbf{M}^T \mathcal{F}_d^{-1} [\mathcal{F}_d [\mathbf{M}\mathbf{H}^T \mathbf{r}] \div (\tilde{\mathbf{c}} + \varepsilon^*)] \quad (\approx \mathbf{A}^{-1}\mathbf{r}) \\ 32 \\ 33 \quad \delta_{k+1} &= \mathbf{r}^T \mathbf{s} \\ 34 \\ 35 \\ 36 \\ 37 \end{aligned}$$

38 3. Update the preconditioned conjugate gradient
39

$$40 \\ 41 \quad \mathbf{d} = \mathbf{s} + \frac{\delta_{k+1}}{\delta_k} \mathbf{d} \\ 42 \\ 43 \\ 44$$

45 4. Increase k by one and repeat until $k > k_{max}$ or $\delta_{k+1} < \varepsilon^2 \delta_0$.
46
47
48
49

50 The variables \mathbf{q} , α and \mathbf{s} are auxiliaries to reduce the computational costs. All products
51 with \mathbf{M} and \mathbf{H} are evaluated by simple embedding/extraction and injection/sampling or
52 using sparse representations of the matrix $\mathbf{N} = \mathbf{M}\mathbf{H}^T$. The PCG algorithm requires only
53 one matrix-vector product per iteration step evaluated via FFT, which has an asymptotic
54 cost estimate of $\mathcal{O}(n \log_2 n)$ and another operation of $\mathcal{O}(n \log_2 n)$ to apply the preconditioner.
55 The corresponding steps in the initialization are evaluated accordingly, resulting
56 in an overall complexity of $\mathcal{O}(n \log_2 n)$. Here, the finer regular grid with n nodes is not
57
58
59
60
61
62
63
64
65

1
2
3 necessarily as large (as fine) as the grid of estimation, subject to the desired accuracy of
4
5 discretizing the locations of the measurements.
6

7
8 In case the measurements themselves lie on a regular grid, matrix $\mathbf{A} = \mathbf{T}$ in the above
9
10 algorithm is a Toeplitz matrix itself, and $\mathbf{H} = \mathbf{H}^T = \mathbf{I}$ may be omitted in the entire
11
12 algorithm. In that case, the computational complexity drops to $\mathcal{O}(m \log_2 m)$.
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

References

- Barel, M. V., Heinig, G., and Kravanja, P. (2001). A stabilized superfast solver for nonsymmetric Toeplitz systems. *SIAM J. Matrix Anal. A.*, 23(2):494–510.
- Barnett, S. (1990). *Matrices Methods and Applications*. Oxford Applied Mathematics and Computing Science Series. Clarendon Press, Oxford.
- Chan, R. H. and Ng, M. K. (1996). Conjugate gradient methods for Toeplitz systems. *SIAM Review*, 38(3):427–482.
- Cirpka, O. A. and Nowak, W. (2003). Dispersion on kriged hydraulic conductivity fields. *Water Resour. Res.*, 39(2):doi:10.1029/2001WR000598.
- Cirpka, O. A. and Nowak, W. (2004). First-order variance of travel time in non-stationary formations. *Water Resour. Res.*, 40:doi:10.1029/2003WR002851.
- Cooley, J. W. and Tukey, J. W. (1965). An algorithm for the machine calculation of complex Fourier series. *Math. Comput.*, 19:297–301.
- Davis, P. J. (1979). *Circulant Matrices*. Pure and Applied Mathematics. John Wiley and Sons, New York.
- Dietrich, C. R. and Newsam, G. N. (1989). A stability analysis of the geostatistical approach to aquifer transmissivity identification. *Stoch. Hydrol. Hydraul.*, 3:293–316.
- Dietrich, C. R. and Newsam, G. N. (1993). A fast and exact method for multidimensional Gaussian stochastic simulations. *Water Resour. Res.*, 29(8):2861–2869.
- Dietrich, C. R. and Newsam, G. N. (1996). A fast and exact method for multidimensional Gaussian stochastic simulations: Extension to realizations conditioned on direct and indirect measurements. *Water Resour. Res.*, 32(6):1643–1652.
- Dietrich, C. R. and Newsam, G. N. (1997). Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. *SIAM J. Sci. Comput.*, 18(4):1088–1107.

- 1
2
3 Duijndam, A. J. W. and Schonewille, M. A. (1999). Nonuniform fast Fourier transform.
4
5 *Geophys.*, 64(2):539–551.
6
7
8 Dykaar, B. B. and Kitanidis, P. K. (1992). Determination of the effective hydraulic
9
10 conductivity for heterogeneous porous media using a numerical spectral approach. 1.
11
12 Method. *Water Resour. Res.*, 28(4):1155–1166.
13
14
15 Fessler, J. A. and Sutton, B. P. (2003). Nonuniform fast Fourier transform using min-max
16
17 interpolation. *IEEE Transactions on Signal Processing*, 51(2):560–574.
18
19
20 Fourmont, K. (2003). Non-equispaced fast Fourier transforms with applications to tomog-
21
22 raphy. *J. Fourier Anal. Appl.*, 9(5):431–450.
23
24
25 Frigo, M. and Johnson, S. G. (1998). FFTW: An adaptive software architecture
26
27 for the FFT. In *Proc. ICASSP*, volume 3, pages 1381–1384, IEEE, Seattle, WA.
28
29 <http://www.fftw.org>.
30
31
32 Gallivan, K., Thirumalai, S., Dooren, P. V., and Vermaut, V. (1996). High performance
33
34 algorithms for Toeplitz and block Toeplitz matrices. *Linear Algebra Appl.*, 241-243(1-
35
36 3):343–88.
37
38
39 Golub, G. H. and van Loan, C. F. (1996). *Matrix Computations*. Jon Hopkins University
40
41 Press, Baltimore, Md, third edition.
42
43
44 Good, I. J. (1950). On the inversion of circulant matrices. *Biometrika*, 37:185–186.
45
46
47 Greengard, L. and Lee, J.-Y. (2004). Accelerating the nonuniform fast Fourier transform.
48
49 *SIAM Review*, 46(3):443–454.
50
51
52 Hestenes, M. R. and Stiefel, E. (1952). Methods of conjugate gradients for solving linear
53
54 systems. *J. Res. Nat. Bur. Stand.*, 49:409–436.
55
56
57 Kailath, T. and Sayed, A. H. (1995). Displacement structure: Theory and applications.
58
59 *SIAM Review*, 37(3):297–386.
60
61
62
63
64
65

- 1
2
3 Kailath, T. and Sayed, A. H. (1999). *Fast Reliable Algorithms of Matrices with Structure*.
4
5 SIAM, Philadelphia, PA.
6
7
8 Kitanidis, P. K. (1993). Generalized covariance functions in estimation. *Math. Geol.*,
9
10 25(5):525–540.
11
12
13 Kitanidis, P. K. (1995). Quasi-linear geostatistical theory for inversing. *Water Resour.*
14
15 *Res.*, 31(10):2411–2419.
16
17
18 Kitanidis, P. K. (1996). Analytical expressions of conditional mean, covariance, and
19
20 sample functions in geostatistics. *Stoch. Hydrol. Hydraul.*, 12:279–294.
21
22
23 Kitanidis, P. K. (1997). *Introduction to Geostatistics*. Cambridge University Press, Cam-
24
25 bridge.
26
27
28 Kozintsev, B. (1999). *Computations with Gaussian random fields*. PhD thesis, Institute
29
30 for Systems Research, University of Maryland.
31
32
33 Liu, Q. H. and Ngyuen, N. (1998). An accurate algorithm for nonuniform fast Fourier
34
35 transforms(NUFFT’s). *IEEE Microwave and Guided Wave Letters*, 8(1):18–20.
36
37
38 Newsam, G. N. and Dietrich, C. R. (1994). Bounds on the size of nonnegative definite
39
40 circulant embeddings of positive definite Toeplitz matrices. *IEEE Transactions on*
41
42 *Information Theory*, 40(4):1218–1220.
43
44
45 Nowak, W. (2005). *Geostatistical Methods for the Identification of Flow and Transport Pa-*
46
47 *rameters in Subsurface Flow*. PhD thesis, Institut für Wasserbau, Universität Stuttgart,
48
49 http://elib.uni-stuttgart.de/opus/frontdoor.php?source_opus=2275.
50
51
52
53 Nowak, W. and Cirpka, O. A. (2004). A modified Levenberg-Marquardt algorithm for
54
55 quasi-linear geostatistical inversing. *Adv. Water Resour.*, 27(7):737–750.
56
57
58 Nowak, W. and Cirpka, O. A. (2006). Geostatistical inference of conductivity and
59
60 dispersion coefficients from hydraulic heads and tracer data. *Water Resour. Res.*,
61
62 42(W08416):doi:10.1029/2005WR004832.
63
64
65

- 1
2
3 Nowak, W., Tenklevé, S., and Cirpka, O. A. (2003). Efficient computation of linearized
4 cross-covariance and auto-covariance matrices of interdependent quantities. *Math.*
5 *Geol.*, 35(1):53–66.
6
7
8
9
10 Pegram, G. G. S. (2004). Spatial interpolation and mapping of rainfall (SIMAR)
11 Vol.3: Data merging for rainfall map production. *Water Research Commission Report*,
12 (1153/1/04).
13
14
15
16
17 Press, W. H., Teukolsky, B. P. F. S. A., and Vetterling, W. T. (1992). *Numerical Recipes:*
18 *The Art of Scientific Computing*. Cambridge University Press, Cambridge, 2nd edition.
19
20
21
22 Rino, C. L. (1970). The inversion of covariance matrices by Finite Fourier Transforms.
23 *IEEE Transactions on Information Theory*, 16:230–232.
24
25
26
27 Schwegge, F. C. (1973). *Uncertain Dynamic Systems*. Prentice-Hall, Englewood Cliffs,
28 NJ.
29
30
31
32 Shewchuk, J. R. (1994). An introduction to the conjugate gradient method without the
33 agonizing pain. <http://www.cs.berkeley.edu/~jrs/>.
34
35
36
37 Strang, G. (1986). A proposal for Toeplitz matrix calculations. *Stud. Appl. Amth.*, 74:171–
38 176.
39
40
41
42 Trapp, G. E. (1973). Inverses of circulant matrices and block circulant matrices. *Kyung-*
43 *pook Math. J.*, 13(1):11–20.
44
45
46
47 van Loan, C. F. (1992). *Computational Frameworks for the Fast Fourier Transform*. SIAM
48 Publications, Philadelphia, PA.
49
50
51
52 Varga, R. S. (1954). Eigenvalues of circulant matrices. *Pacific J. Math.*, 4:151–160.
53
54
55 Wesson, S. M. and Pegram, G. G. S. (2004). Radar rainfall image repair techniques.
56 *Hydrological and Earth Systems Sciences*, 8(2):8220–234.
57
58
59
60 Zimmerman, D. L. (1989). Computationally exploitable structure of covariance matrices
61 and generalized covariance matrices in spatial models. *J. Stat. Comput. Sim.*, 32(1/2):1–
62 15.
63
64
65

Figures

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 1:

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

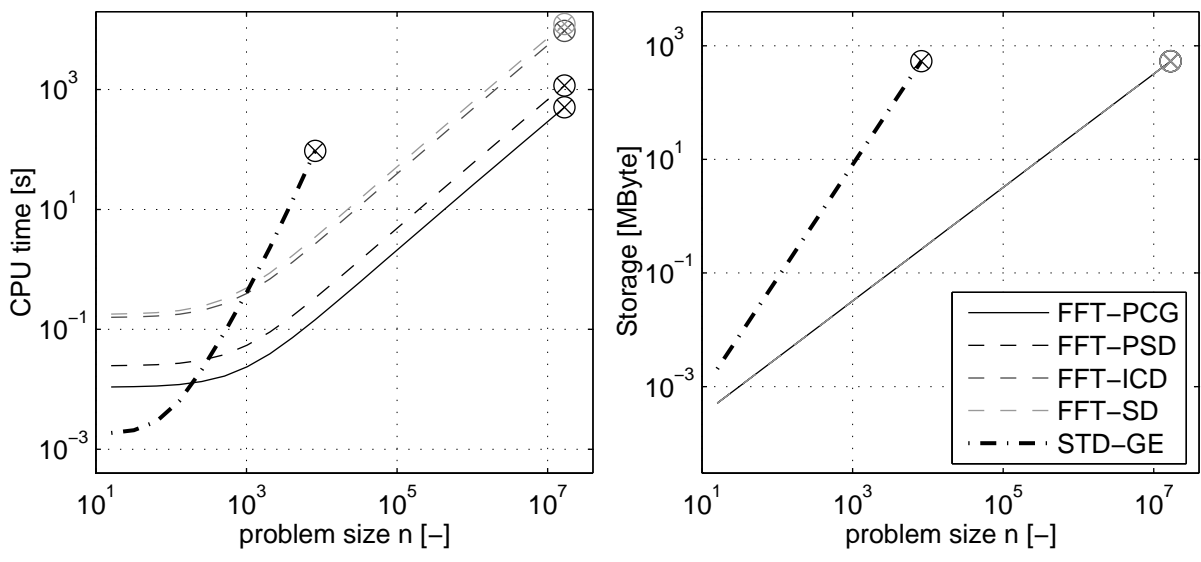


Figure 2:

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

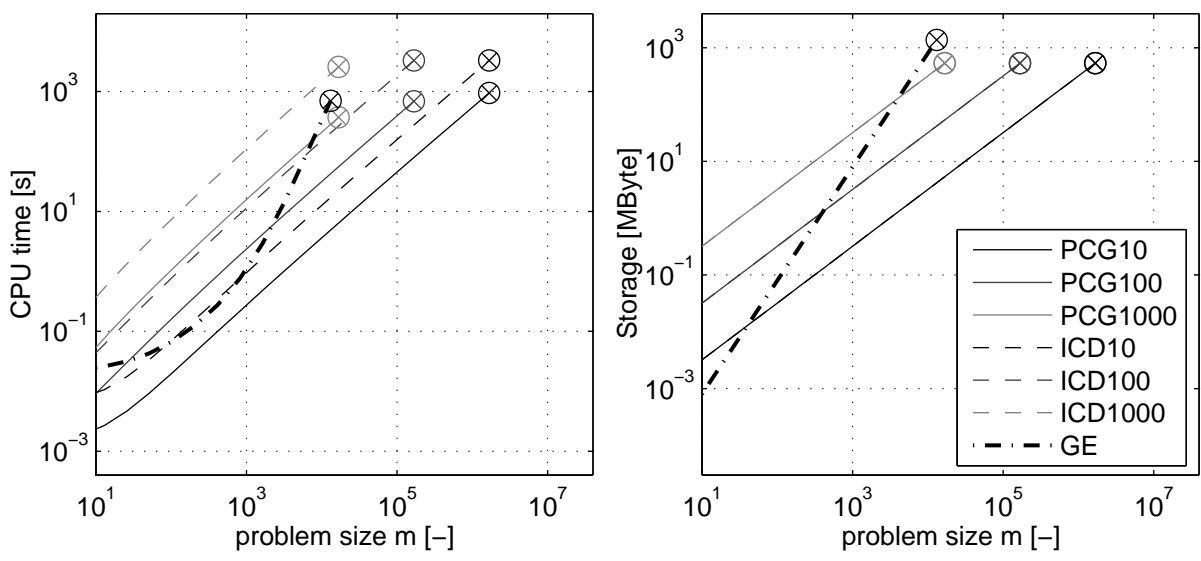


Figure 3:

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 4:

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 5:

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 6:

Captions:

Figure 1: Convolution/Superposition: CPU time (left) and storage requirements (right) for conventional discrete superposition versus superposition via FFT for different problem sizes n and numbers of superimposed terms m . Lines: fitted complexity models. X-Circles: memory overflow. Solid line: superposition/convolution via FFT (same for both tasks). Dashed lines: superposition via brute-force matrix product at different ratios n/m . Dash-dotted line: convolution via matrix product($n = m$). Lines exceeding point of memory overflow: superposition/convolution via successive addition of shifted kernel function. Lower line ends: minimum of $m = 1$ superimposed term.

Figure 2: Solvers for regular measurement grids: CPU time (left) and storage requirements (right) of different solution techniques for Toeplitz systems for different problem sizes m . Lines: fitted complexity models. X-Circles: memory overflow. FFT-PCG: FFT-based PCG solver. FFT-PSD: FFT-based Preconditioned Steepest Descent. FFT-ICD: Iterative Constrained Deconvolution. FFT-SD: FFT-based Steepest Descent. STD-GE: Standard Gaussian Elimination.

Figure 3: Extension to irregular measurement grids: CPU time (left) and storage requirements (right) of different solution techniques for nearly Toeplitz systems for different problem sizes m and different sizes n of the embedding Toeplitz matrix. Lines: fitted complexity models. X-Circles: memory overflow. PCG#: FFT-based PCG solver with # times larger regular grid. ICD#: FFT-based Iterative Constrained Deconvolution with # times larger regular grid. GE: Standard Gaussian Elimination.

Figure 4: Performance of Kriging with conventional solver: CPU time (left) and storage requirements (right) for different numbers n of estimation points and different numbers m of measurements. Lines: fitted complexity models. X-Circles: memory overflow. Solid lines: using conventional solver and conventional superposition. Dashed lines: using conventional solver and superposition via FFT.

Figure 5: Performance of Kriging with FFT-based PCG solver for irregularly scattered

1
2
3 measurements: CPU time (left) and storage requirements (right) for different numbers n
4
5 of estimation points and different numbers m of measurements. Lines: fitted complexity
6
7 models. X-Circles: memory overflow. Solid lines: exact estimation variance. Dashed
8
9 lines: one-point approximation.
10

11
12 Figure 6: Performance of Kriging with FFT-based PCG-solver with measurements on
13 regular grids: CPU time (left) and storage requirements (right) for different numbers n
14
15 of estimation points and different numbers m of measurements. Lines: fitted complexity
16
17 models. X-Circles: memory overflow. Solid lines: exact estimation variance. Dashed
18
19 lines: large-grid approximation.
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65