Next, we demonstrate phase retrieval from partial measurements, that is, combining phase retrieval with super-resolution. Conceptually, this experiment corresponds to a discrete version of recent work on sparsity-based sub-wavelength coherent diffractive imaging [6,7].

Figure 4 presents the results of experiments showing the recovery of a 3-sparse signal (Figs. 4(b) and 4(d) show the signal's amplitude and phase, respectively), from one half of the available measurements (optical intensities at the output of the waveguides), obtained by discarding every other measurement. Figure 4(e) shows the used part of the discretized measurements. The measured point-spread-function of the system is shown in Fig. 4(a), and the measured output intensities before discretization is shown in Fig. 4(c). Figure 4(f) demonstrates the super-resolution effect achieved by the method – where the non-measured intensities (every other waveguide) are in fact recovered with high accuracy. Here, waveguide 13 (the left input waveguide) is chosen as having phase = 0. The correspondence between the amplitude and phase of the recovered signal and the amplitude and phase of original input, displayed in Figs. 4(b), 4(d), demonstrates that our sparsity-based technique performs super-resolution with phase retrieval accurately.
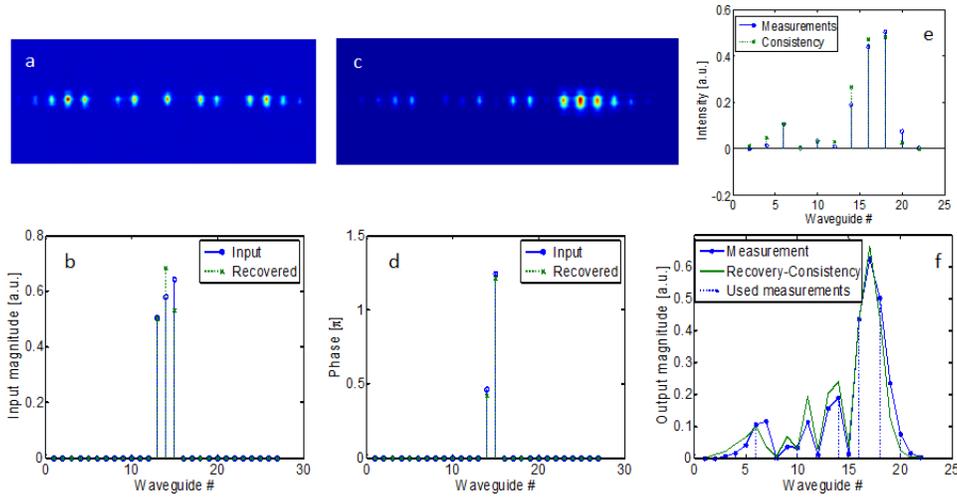

Fig. 4. Recovery from partial measurements. (a) Point spread function of the system. (b) Input and recovered amplitudes. (c) Measured output intensities. (d) Input and recovered phases. (e) The 11 measurements used and the consistency of the recovered signal at the system's output. (f) Super-resolution: accurate recovery of the non-measured output intensities (blue) using half of the output intensities (dashed blue).

Numerical simulations

Having demonstrated the power of our sparsity-based method in phase-retrieval combined with super-resolution, it is important to study the generality of the method in terms of sparsity and noise. To do that, we perform numerical simulations, displayed in Fig. 5(a). In the simulations, we generate random complex input signals, which yield their corresponding partial output intensities using the transfer function of Eq. (1). We use GESPAR to recover the input signal from part of the measurements vector. Figure 5(a) shows the normalized reconstruction error, defined as $\dfrac{\left\| x - x_{rec} \right\|_2}{\left\| x \right\|_2}$, vs. the input signal's sparsity level k, for different SNR levels. In these simulations, one half of the intensity measurements are used for reconstruction (20 out of 40), obtained by keeping only every other output waveguide intensity information. Each input signal is composed of k non-zero elements at randomly chosen locations with uniformly random phase in the range $[0, 2\pi]$, and uniformly distributed

magnitude in the range $[1,2]$. Each data point is the averaged result of 100 iterations. Figure 5(a) shows that good signal recovery (amplitude and phase recovery with error comparable to the noise level) of up to a sparsity level of 4. That is, our method facilitates super-resolution combined with phase retrieval when up to 4 waveguides are excited out of the possible 21 waveguides at the center of our array, given 20 intensity measurements and SNR of 30 (corresponding to ~3% of noise) and higher.

Another aspect we investigate numerically is the number of measurements necessary to recover a signal robustly for different sparsity levels. Figure 5(b) shows the probability to correctly recover the excited waveguides (the support of the signal), as a function of signal sparsity, for various levels of m – the number of measurements. The measurements are distributed evenly in the output plane, and the SNR is 40. As can be seen, the support recovery probability is very high (>95%) for sparsity of 4, using 24 measurements, out of the possible 40 over which the discrete diffraction pattern can extend given the propagation distance z and coupling constant c, which were taken such that $cz = 4$. If the sparsity is very low ($k = 2$), 16 measurements empirically guarantee a robust recovery under the conditions of the simulation.
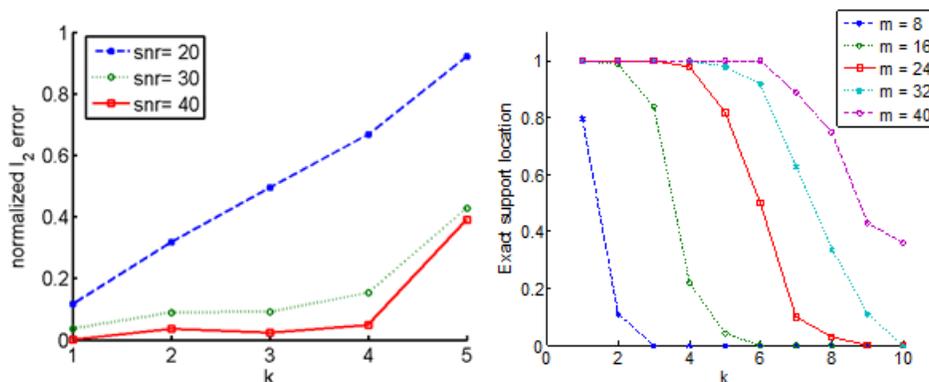

Fig. 5. (a) Reconstruction error vs. sparsity level (k) for various SNR values. (b) Support location vs. sparsity levels for different values of m - number of evenly spaced intensity measurements

Discussion

An interesting aspect to investigate is the uniqueness of the system, namely, under what circumstances does a set of measurements correspond to a unique input vector. More practically, in the presence of noise – when is it guaranteed that the recovered input vector is close to the true input? Without using the sparsity information, the system is very ill posed, and multiple, significantly different, input signals produce very similar output intensities. This can be easily verified by using GESPAR with a large sparsity parameter s and obtaining arbitrary non-sparse signals that are local minima of the problem in Eq. (4). However, as the simulations indicate, using the sparsity prior regularizes the inverse problem, allowing robust recovery of sparse signals with high probability.

In this context, we note that, while theoretical guarantees of robust sparse recovery from quadratic measurements do exist in the context of random measurements [16], the measurements in our system are structured, and therefore these guarantees do not directly apply. In other related recent work [20], uniqueness guarantees and recovery robustness of general phase-retrieval problems have been studied - but for general (non-sparse) signals, yielding sufficient uniqueness conditions on the measurement matrix, that are very hard to verify. However, in our case, the additional prior knowledge of signal sparsity may significantly improve these guarantees. Robustness guarantees for structured measurements of sparse signals are the subject of current work in progress, see for example [21].

One interesting property that the coupled waveguide system possesses is the following specific phase ambiguity: Denoting a certain excited waveguide at the input as waveguide 0 and setting its arbitrary phase to 0, then an identical output intensity will be attained by transforming the phase of all input waveguides according to: $\phi_k \to k\pi - \phi_k$ where k is the waveguide index. This ambiguity is a mathematical property of the transfer function, and as such it is inherent in the system, similar to the mirroring and shifting ambiguities occurring in the Fourier phase retrieval problem. This ambiguity exists here even in the completely noiseless case.

Another aspect to be examined is the application of the scheme as a communication device. For example, when will it be beneficial to use the suggested scheme over simply using m uncoupled waveguides? On one hand, limiting the input signal to be sparse decreases the number of different possible transmittable signals. On the other hand, only by using coupling, one can recover the input phase information (In this sense the system can be seen as an interference measurement), so that a new degree of freedom to encode information is introduced. Whether using the sparse mode coupled system is beneficial will depend on the specific properties of the system, namely signal to noise ratio (SNR), coupling coefficient, etc., but a general framework of comparing certain aspects of the two approaches is presented.

To make a fair comparison, we allow multilevel intensity detection in both cases. Let us first consider a system of m uncoupled waveguides, with m intensity measurements. Multilevel detection implies non binary quantization of the input signal, and we denote by q the number of quantization levels that would allow a robust recovery (up to a pre-determined bit error rate) given the system's SNR. The number of different transmittable signals in the system is therefore $n_{uc} = q^m$.

Moving on to the coupled-waveguide system – we denote by $q_a{'}$ and $q_p{'}$ the number of quantization levels in a k sparse signal's amplitude and phase, respectively, that would allow its recovery using an array of n waveguides, given m intensity measurements . The number of possible signals is in this case is given by:

$$n_c = q_a{'}^k \cdot q_p{'}^k \cdot \binom{m}{k} = \frac{q_a{'}^k \cdot q_p{'}^k \cdot m!}{k!(m-k)!}$$

(5)

Using the coupled mode system will therefore be beneficial in the sense of the number of possible transmittable signals if:

$$\alpha = \frac{n_c}{n_{uc}} = \frac{q_a{'}^k \cdot q_p{'}^k \cdot m!}{q^m \cdot k! \cdot (m-k)!} > 1$$

(6)

The above discussion is only partial, and many other considerations apply to an actual communication system, e.g. the process necessary to recover the input given the output. For example, in an actual communication application, the signal recovery would have to be very fast in order for the scheme to be useful. The exact performance details will depend on the implementation and on signal size and sparsity level, however, GESPAR has been shown to outperform existing phase retrieval algorithms in terms of computational speed as well as recovery performance [9]. For small enough signal sizes and sparsity the algorithm is very fast, requiring a small number of matrix inversions, the size of which is linear in the number of waveguides times the sparsity level.

Finally, it should be noted that having additional prior information on the input signal can improve the recovery performance. Adding prior information in addition to sparsity in order to improve recovery performance has been recently suggested in the context of image interpolation [22]. In our case, such additional information may be the total power in the system, a known amplitude relation between the input waveguides (e.g. all input waveguides have identical input amplitude - a binary input signal), or known phase relation - e.g. all input

waveguides have the same input phase. The latter can be easily incorporated into GESPAR, as it can be viewed as a real (rather than complex) input signal, which is the case dealt with in [9].

Conclusion

In conclusion, we have demonstrated numerically and experimentally super-resolution phase-retrieval in an array of coupled waveguides. This is a generalization of sparsity based super resolution combined with phase retrieval [6] to a system with a transfer function different than free space propagation. This waveguide array serves as a model for a general system of optical interconnects, characterized by a specific transfer function. The work presented here has studied a one dimensional array, but naturally, interconnects can also have a higher dimensionality – depending on the number of nearest neighbors coupled to every specific site. From the recent paper on super-resolution coherent diffractive imaging, we conjecture that two dimensional interconnects would most probably also allow sparsity-based super-resolution phase-retrieval. It is most certainly interesting to explore phase retrieval in interconnects of a higher dimensionality. In a more general context, the method proposed here can be applied to various other multiple-input multiple-output (MIMO) communication schemes – for instance schemes based on multimode or multi-core fibers.

In addition, the technique can serve as the basis for super-resolution in quantum coincidence measurements in a waveguide array setting [23–25]. The ability to efficiently measure super-positions of states consisting of several photons is essential the characterization of quantum optical systems and computation units [26,27]. Such measurements are typically coincidence measurements. However, high-order coincidence measurements, which are required for the characterization of systems of more than two photons, are very hard to implement experimentally, and the number of different measurements required for the n-fold coincidence grows very fast with n. Sparsity may be used to reduce the number and the dimension of the required coincidence measurements, and a coupled waveguide array system can be suitable for an experimental demonstration of this concept [28].