SI Text

Fluorescence Measurements. The QDs are suspended as a dilute solution in a 9:1 mixture of hexane and octane, and a drop of QD solution is evaporated onto a glass microscope coverslip. The concentration of the solution is adjusted to give well isolated single dots on the surface. The immobilized QDs are illuminated with light from a low-noise, 532-nm laser (Spectra-Physics Excelsior), which is attenuated to various intensities by using neutral-density filters. The incident power is adjusted in the range of 2.5-7.5 μW, corresponding to a focused intensity of ≈4-12 kW/cm^2. This relatively high intensity was used to obtain a high photon count rate, enabling the subsequent statistical analysis. The excitation light is focused onto a single QD by using an oil-immersion microscope objective (Olympus PLAN APO ×60); the same objective is used to collect QD fluorescence. The collected emission is separated from scattered laser light by using a dichroic mirror (Chroma 565 DCXR), and is further isolated by using a band-pass filter (Chroma HQ615/60m). It is split into two equal parts by using a beamsplitter, and the two beams are coupled into optical fibers and detected by using single-photon-counting avalanche photodiodes (Perkin-Elmer model SPC-AQR-15). Pulses from the detectors are fed through a router (Becker & Hickl HRT-82) and sent to a data-acquisition board (Becker & Hickl SPCM-600). The board is configured in time-tag (FIFO48) mode, using a 40-MHz sync signal, time-to-amplitude conversion (TAC) range of 50 nsec, TAC offset of 5%, analog-to-digital conversion resolution of 4,096, and dither range of 1/8. The TAC is started by a detected photon and stopped by the next sync pulse. This configuration records the arrival time of each photon since the start of the experiment with 50-nsec resolution. The measured background photon detection rate, due to dark counts, scattered laser light, and background fluorescence, is ≈100 counts per sec.

Autocorrelation Functions. If the autocorrelation function is calculated from the photon count rate at a single detector, it will be dominated at short time delays by the effects of detector afterpulsing: shortly after a real photon-detection event, the detector has a probability of recording a second, false, event. To avoid this artifact, the fluorescence light is split into two equal parts, which are sent to separate detectors. $G^{(2)}(\tau)$ is then approximated as the cross-correlation function between the photon count rates, $i_1(t)$ and $i_2(t)$, at the two detectors. Because the photon count rates are measured as discrete time series $\{i_1(t_1), i_1(t_2), i_1(t_3), \ldots, i_1(t_N)\}$ and $\{i_2(t_1), i_2(t_2), i_2(t_3), \ldots, i_2(t_N)\}$, the discrete-time autocorrelation function is calculated from these series as

$$G^{(2)}(\tau_m) = \frac{1}{N - |m|} \sum_{n=0}^{N-m-1} i_1(t_n) i_2(t_{n+m}). \quad \textbf{[S1]}$$

This "unbiased" estimate of the autocorrelation avoids rolloff, due to the finite duration of the experiment, that can otherwise dominate the form of $G^{(2)}(\tau)$ at longer τ.

The autocorrelation function is determined by breaking the full time experimental time series into blocks of ≈10 sec, calculating $G^{(2)}(\tau)$ for each block, and averaging together the results to give the final autocorrelation. This procedure greatly reduces computational time, at the expense of reducing the maximum time delay τ that is investigated. Calculated correlation functions are rebinned logarithmically: the values of τ are divided into a series of intervals, equally spaced along a logarithmic axis, and the average value of the correlation function is calculated in each of these intervals. Background counts due to stray photons and dark counts are expected to be uncorrelated both to one another and to photons emitted by the QD, and will not contribute to the autocorrelation function for finite value of τ. The calculated $G^{(2)}(\tau)$ thus reflects the underlying statistical properties of photon emission by the QD, including blinking processes.

Power Spectra. To avoid the effects of afterpulsing in the single-photon detectors, we split the photons and send them toward two detectors, as for the determination of the autocorrelation function. The power spectrum is then approximated as $P(f) \approx I_1(f) I_2^*(f)$, where $I_1(f)$ and $I_2(f)$ are the Fourier transforms of $i_1(t)$ and $i_2(t)$, respectively. In the limit of an infinite measurement time, this approximation will converge on the true power spectrum; for real data, the approximation returns some small imaginary and negative components, which are neglected.

Long time series are broken into shorter blocks, as for the autocorrelation functions. Before the blocks are Fourier transformed, they are multiplied by a Blackmann-Harris window, to minimize numerical artifacts at the edges of the frequency range. The resulting power spectra are rebinned logarithmically in frequency. Even with the application of the window function, the lowest-frequency point is dominated by the effects of the finite length of the blocks, and is therefore eliminated. This entire computational procedure was tested on simulated data streams, verifying that it provides an accurate representation of the underlying power spectral density. We have also calculated the numerical inverse Fourier transform of the calculated power spectra, which should be equal to the directly calculated autocorrelation functions. The inverse Fourier transform produced small numerical artifacts at long time delays, due to the finite frequency bandwidth of the power spectra; otherwise, good agreement was observed.

The power spectrum, like the autocorrelation function, is calculated from a finite time series, and is thus an approximation of the ideal power spectrum that would result from an experiment of infinite duration. The statistical error, reflecting the expected difference between the calculated and the ideal power spectra, is reduced by the practice of dividing the time series into *K* shorter blocks and averaging the calculated spectra for each of those blocks. The error is further reduced by rebinning the averaged spectrum logarithmically in frequency; it is this rebinned spectrum that is shown in Fig. 4. The resulting standard deviation at frequency f is

$$\sigma(f) = \frac{P(f)}{\sqrt{K n_{bin}(f)}}, \qquad \textbf{[S2]}$$

where $n_{bin}(f)$ is the number of frequency points that are rebinned together at frequency f. Evaluating this error at each point in Fig. 4 yields a statistical error that is smaller than the size of the points on the graph.

The power spectrum includes the contribution of both “on” and “off” blinking processes and will be dominated by the process that contributes the greater noise power. If the two processes are described by different power laws, then there may be a time scale below which one process dominates and above which the other process dominates. This would also produce a change in the slope of the power spectrum; in this case, however, the magnitude of the slope would be greater at low frequencies than at high frequencies, whereas we observe the opposite change in slope. More specifically, if the total power spectrum were the simple combination of power laws due to on and off blinking statistics, then it would have the form

$$P(f) = C_{on} f^{-m_{on}} + C_{off} f^{-m_{off}}, \qquad \textbf{[S3]}$$

where C_{on}, C_{off}, $m_{on} \equiv -\nu_{on} + 2$, and $m_{off} \equiv -\nu_{off} + 2$ are all positive. It is then straightforward to see that

$$\frac{dP(f)}{df} = -C_{on} m_{on} f^{-m_{on}-1} - C_{off} m_{off} f^{-m_{off}-1} < 0, \qquad \textbf{[S4]}$$

and

$$\frac{d^2P(f)}{df^2} = C_{on} m_{on} (m_{on} + 1) f^{-m_{on}-2} + C_{off} m_{off} (m_{off} + 1) f^{-m_{off}-2} > 0, \quad \textbf{[S5]}$$

so that the total power spectrum would be concave. This is illustrated schematically in Fig. 6. What we observe, on the other hand, is a convex power spectrum. Thus, the probability density of either the bright periods or the dark periods, or both, must be changing to produce the observed change in the power spectrum.