

Very-high-frequency oscillations in the main peak of a magnetar giant flare

<https://doi.org/10.1038/s41586-021-04101-1>

Received: 17 August 2020

Accepted: 6 October 2021

Published online: 22 December 2021

 Check for updates

A. J. Castro-Tirado^{1,2}, N. Østgaard³, E. Göğüş⁴, C. Sánchez-Gil⁵, J. Pascual-Granado¹, V. Reglero^{6,7}, A. Mezentsev³, M. Gabler⁶, M. Marisaldi^{3,8}, T. Neubert⁹, C. Budtz-Jørgensen⁹, A. Lindanger³, D. Sarria³, I. Kuvvetli⁹, P. Cerdá-Durán⁶, J. Navarro-González⁷, J. A. Font^{6,10}, B.-B. Zhang^{11,12,13}, N. Lund⁹, C. A. Oxborrow⁹, S. Brandt⁹, M. D. Caballero-García¹, I. M. Carrasco-García¹⁴, A. Castellón^{2,15}, M. A. Castro Tirado^{1,16}, F. Christiansen⁹, C. J. Eyles⁷, E. Fernández-García¹, G. Genov³, S. Guziy^{17,18}, Y.-D. Hu^{1,19}, A. Nicuesa Guelbenzu²⁰, S. B. Pandey²¹, Z.-K. Peng^{11,12}, C. Pérez del Pulgar², A. J. Reina Terol², E. Rodríguez¹, R. Sánchez-Ramírez²², T. Sun^{1,23,24}, K. Ullaland³ & S. Yang³

Magnetars are strongly magnetized, isolated neutron stars^{1–3} with magnetic fields up to around 10^{15} gauss, luminosities of approximately 10^{31} – 10^{36} ergs per second and rotation periods of about 0.3–12.0 s. Very energetic giant flares from galactic magnetars (peak luminosities of 10^{44} – 10^{47} ergs per second, lasting approximately 0.1 s) have been detected in hard X-rays and soft γ -rays⁴, and only one has been detected from outside our galaxy⁵. During such giant flares, quasi-periodic oscillations (QPOs) with low (less than 150 hertz) and high (greater than 500 hertz) frequencies have been observed^{6–9}, but their statistical significance has been questioned¹⁰. High-frequency QPOs have been seen only during the tail phase of the flare⁹. Here we report the observation of two broad QPOs at approximately 2,132 hertz and 4,250 hertz in the main peak of a giant γ -ray flare¹¹ in the direction of the NGC 253 galaxy^{12–17}, disappearing after 3.5 milliseconds. The flare was detected on 15 April 2020 by the Atmosphere–Space Interactions Monitor instrument^{18,19} aboard the International Space Station, which was the only instrument that recorded the main burst phase (0.8–3.2 milliseconds) in the full energy range (50×10^3 to 40×10^6 electronvolts) without suffering from saturation effects such as deadtime and pile-up. Along with sudden spectral variations, these extremely high-frequency oscillations in the burst peak are a crucial component that will aid our understanding of magnetar giant flares.

We report here the detection¹¹ of a new giant flare (initially dubbed GRB 200415) with the Atmosphere–Space Interactions Monitor (ASIM) aboard the International Space Station (ISS) on 15 April 2020 at 08:48:05.56 (± 0.03) UT. With the ASIM Modular X- and Gamma-ray Sensor (MXGS) instrument^{18,19}, we recorded data for 2 s centred around the burst. The two independent detectors of MXGS, covering energies between 50–400 keV (low-energy detector, LED) and 300 keV to 40 MeV (high-energy detector, HED), did not suffer from saturation effects (deadtime, pile-up) and recorded for the first time the fine structure of the main burst phase (0.8–3.2 ms) of a magnetar in this entire energy

range (Fig. 1, Extended Data Fig. 2). Owing to the large effective area of ASIM, 1- μ s time resolution and large energy range, we have performed both a detailed time analysis and a spectral analysis of the main phases of the giant flare. We are able to resolve the complex temporal structure prior to the absolute peak emission, consisting of six distinct intensity peaks during the first 3.2 ms, this flare being the first one for which we have seen multiple peaks prior to the maximum (see Fig. 1). During the approximately 160-ms duration of the giant flare, around 10^{46} erg isotropic equivalent energy was released, roughly the energy the Sun radiates in about 100,000 yr.

¹Instituto de Astrofísica de Andalucía del Consejo Superior de Investigaciones Científicas (IAA-CSIC), Granada, Spain. ²Unidad Asociada al CSIC Departamento de Ingeniería de Sistemas y Automática, Escuela de Ingeniería Industrial, Universidad de Málaga, Málaga, Spain. ³Birkeland Centre for Space Science, Department of Physics and Technology, University of Bergen, Bergen, Norway. ⁴Faculty of Engineering and Natural Sciences, Sabanci University, Istanbul, Turkey. ⁵Departamento de Estadística e Investigación Operativa, Universidad de Cádiz, Puerto Real, Spain. ⁶Departamento de Astronomía y Astrofísica, Universitat de València, Valencia, Spain. ⁷Imaging Processing Laboratory, Universitat de València, Valencia, Spain. ⁸National Institute for Astrophysics, Osservatorio di Astrofisica e Scienza dello Spazio, Bologna, Italy. ⁹DTU Space, Technical University of Denmark, Kongens Lyngby, Denmark. ¹⁰Osservatori Astronomic, Universitat de València, Valencia, Spain. ¹¹School of Astronomy and Space Science, Nanjing University, Nanjing, China. ¹²Key Laboratory of Modern Astronomy and Astrophysics, Nanjing University, Ministry of Education, Nanjing, China. ¹³Department of Physics and Astronomy, University of Nevada, Las Vegas, Las Vegas, NV, USA. ¹⁴Departamento de Química Analítica, Universidad de Málaga, Málaga, Spain. ¹⁵Departamento de Álgebra, Geometría y Topología, Facultad de Ciencias, Universidad de Málaga, Málaga, Spain. ¹⁶Escuela Técnica Superior de Arquitectura, Universidad de Málaga, Málaga, Spain. ¹⁷Astronomical Observatory, Mykolaiv National University, Mykolaiv, Ukraine. ¹⁸Research Institute, Mykolaiv Astronomical Observatory, Mykolaiv, Ukraine. ¹⁹Facultad de Ciencias, Universidad de Granada, Granada, Spain. ²⁰Thüringer Landessternwarte Tautenburg, Tautenburg, Germany. ²¹Aryabhata Research Institute of Observational Sciences (ARIES), Nainital, India. ²²Istituto di Astrofisica e Planetologia Spaziali, INAF, Rome, Italy. ²³Purple Mountain Observatory, Chinese Academy of Sciences, Nanjing, China. ²⁴School of Astronomy and Space Science, University of Science and Technology of China, Hefei, China. ✉e-mail: Nikolai.Ostgaard@uib.no; ersin.gogus@sabanciuniv.edu; Andrey.Mezentsev@uib.no; michael.gabler@uv.es; Martino.Marisaldi@uib.no

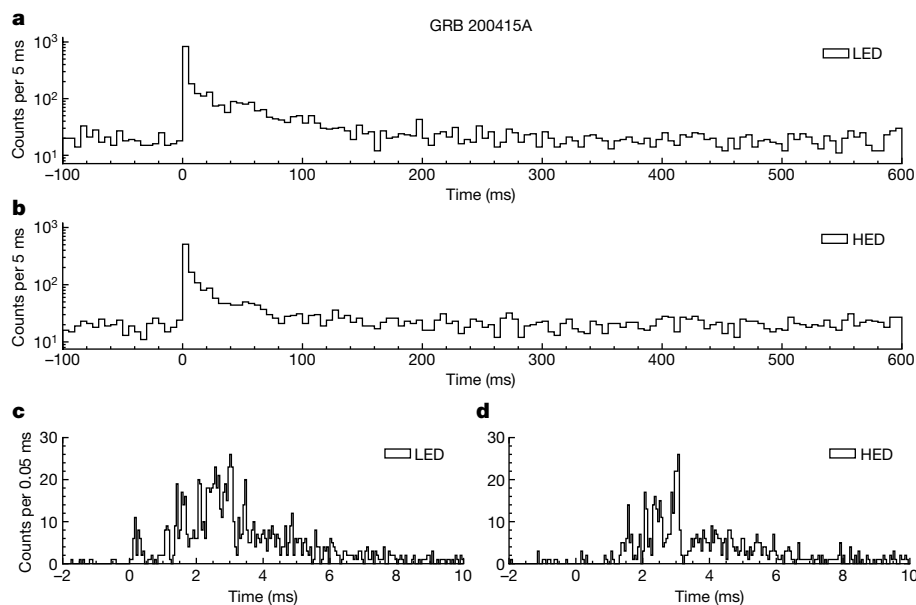


Fig. 1 | Temporal variability of GRB 200415A. **a**, The ASIM LED (low-energy detector) hard X-ray light curve of the magnetar giant flare. **b**, The ASIM HED (high-energy detector) γ -ray light curves of the magnetar giant flare. Both light curves have a

complex structure, displaying six intensity peaks during the first 3.2 ms. The time interval ranges from -100 to $+600$ ms, with the data binned to 5 ms.

c, d, A magnification of the first 10 ms with the LED (**c**) and HED (**d**) data binned to 50 μ s.

On the basis of the temporal structure and energy distribution we distinguish four distinct phases: (i) precursor (0–0.8 ms, with flux dominated by low-energy counts); (ii) burst (0.8–3.2 ms, with well defined peak structure); (iii) decay (3.2–8.0 ms); and (iv) tail (8–160 ms). For all four distinct phases we created detailed energy spectral fits (see Methods for details), which are summarized in Fig. 2, Table 1 and Methods (Extended Data Fig. 3). The preferred values for all of the four phases are marked with bold text in Table 1.

The spectrum of the precursor phase is dominated by a blackbody with a temperature of 66 keV. During the peak interval (0.8–3.2 ms) the flux increases by one order of magnitude and we see a major non-thermal component that can be modelled by a power law with exponential cut-off with peak energy of 1,160 keV (maximum of the spectral energy distribution). The flux of this non-thermal component then gradually decreases during the decay and tail phase. In the tail phase the blackbody component reappears, now with a slightly lower temperature (40 keV) than in the precursor phase.

To search for quasi-periodic oscillations in the signal, we created individual light curves each with a time resolution of 50 μ s for the LED and HED data, with four different search windows: 0–5 ms, 0–10 ms, 0–50 ms and 0–100 ms. The resulting power spectral density (PSD) reaches a Nyquist frequency of 10^4 Hz, which is high enough to include the shortest dynamical timescales for a typical neutron star. The Leahy-normalized²⁰ PSD constructed using the LED data reveals two statistically significant QPOs at frequencies $2,156 \pm 150$ Hz and

$4,256 \pm 323$ Hz (Fig. 3a, Extended Data Fig. 5, Extended Data Table 2, marked with bold). These oscillations are seen exclusively during the first few milliseconds, which includes the main burst phase (0.8–3.2 ms; see Fig. 1c, Fig. 3b, and Methods for the details of the PSD modelling), and they vanish 3.5 ms after the burst onset (Extended Data Fig. 8c). The detection of the QPO coincides with the increase of the non-thermal emission. Because 4,256 Hz is twice that of the lower frequency oscillation within errors, it is probably the first harmonic of the broad feature at 2,156 Hz. The signal-to-noise ratio (SNR) at 2,156 Hz corresponds to a P value $< 10^{-3}$ (see Extended Data Table 2 for additional details). The PSD constructed using the HED data also shows these QPOs at $2,185 \pm 52$ Hz and $4,200 \pm 37$ Hz (Extended Data Table 2, Extended Data Fig. 6), and both are in agreement with the frequencies from the LED data within errors. Additionally, the HED power spectrum indicates the presence of another statistically significant oscillation at approximately $1,487 \pm 33$ Hz, marked with bold in Extended Data Table 2.

We have also searched the unbinned LED and HED event arrival times for these signals using Z_m^2 statistics, where m is the number of harmonics over which the Fourier powers are added²¹. We selected $m = 1$, for which Z^2 values correspond to twice the powers of the Rayleigh test. We applied this technique in a dynamic manner. In particular, we considered a search interval of 5 ms, moving in steps of 1 ms with the start time ranging from -6 ms through to 50 ms with respect to the burst trigger time. In the frequency domain, we searched in 1,000 steps within a 600-Hz frequency interval centred on the above-reported oscillation

Table 1 | Best-fitting parameters for the spectral properties of the different time epochs

Phase	Time interval (ms)	Model ^a	Photon index	Peak energy (keV)	Temperature (keV)	Reduced χ^2 (d.o.f.)	Flux ^d ($\times 10^{-5}$ erg cm ⁻² s ⁻¹)
Precursor	0–0.8	CPL	$0.4^{+1.5}_{-1.1}$	281^{+84}_{-59}	- ^b	0.50 (11) ^c	$10.2^{+5.6}_{-3.6}$
		BB	- ^b	- ^b	66^{+15}_{-12}	0.52 (12)^c	$9.7^{+12}_{-5.4}$
Peak	0.8–3.2	CPL	$-0.89^{+0.12}_{-0.11}$	$1,160^{+120}_{-100}$	- ^b	1.49 (24)	$85.7^{+7.7}_{-7.1}$
Decay	3.2–8.0	CPL	$-0.41^{+0.16}_{-0.15}$	$1,400^{+110}_{-100}$	- ^b	1.01 (24)	$44.2^{+5.1}_{-4.6}$
Tail	8.0–160	CPL	$-0.85^{+0.15}_{-0.14}$	990^{+140}_{-110}	- ^b	1.96 (21)	$2.26^{+0.16}_{-0.15}$
		BB+CPL	$6.3^{+3.0}_{-1.0}$	737^{+34}_{-32}	$39.6^{+3.3}_{-3.0}$	0.71 (19)	$2.19^{+0.20}_{-0.19}$

^aCPL, cut-off power law; BB, blackbody.

^cC statistics are used instead of χ^2 owing to the low number of counts available.

^dIn the 60 keV–10 MeV range, 68% confidence.

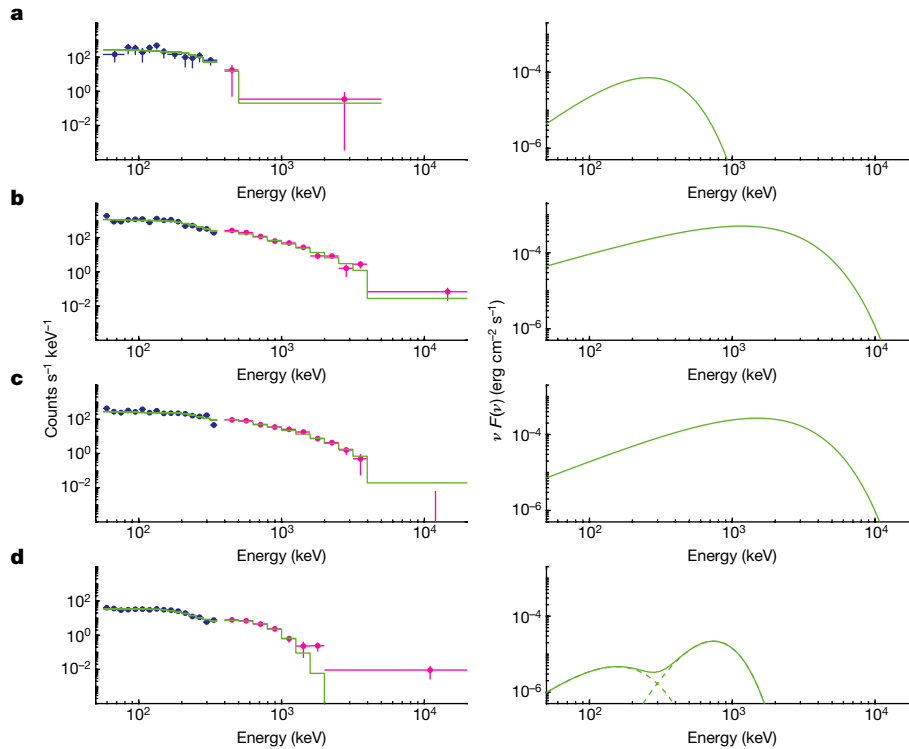


Fig. 2 | Time-resolved spectroscopy. Spectra for four time intervals during the magnetar giant flare, defined in the text. **a**, Precursor: 0.0–0.8 ms. **b**, Peak: 0.8–3.2 ms. **c**, Decay: 3.2–8 ms. **d**, Tail: 8–160 ms. A total of 2,148 counts in the

LED and 2,141 counts in the HED have been used. Left, count spectra (blue, LED; red, HED) and best-fit model (green). Right, energy flux density for the best-fit models listed in Table 1.

frequencies. For the LED, we obtained the highest Z^2 power at $f_1 = 2,132$ Hz (Extended Data Table 3) in the time interval from 0 to 5 ms (see Fig. 3b), both consistent with the PSD results. We determined the chance probability of this signal as 2.4×10^{-9} (Extended Data Table 3). A Z_1^2 search at even higher frequencies yields a peak power at $f_2 = 4,250$ Hz with a chance probability of 1.7×10^{-4} in the same time interval. This strengthens the evidence that indicates that the highest frequency signal is most probably the harmonic. Similarly, we have performed Z_1^2 searches in exactly the same manner using the HED data. We obtain clear signals at 2,095 Hz with a chance probability of 5.0×10^{-8} , and at 4,127 Hz with a probability of arising from a random data sample as 1.1×10^{-2} . We also conducted another search centred on 1,400 Hz, evident from the PSD modelling of the HED data. We find a significant Z_1^2 peak at 1,353 Hz with a chance probability of 1.2×10^{-12} in HED but much less significant (probability of approximately 0.05) in LED (see Extended Data Table 3). This Article analyses only the QPOs significant in both HED and LED, and in the following we will refer to these frequencies, $f_1 = 2,132$ Hz and $f_2 = 4,250$ Hz, as identified in LED data.

We have also investigated unweighted Swift/Burst Alert Telescope (BAT) Gamma-ray Urgent Archiver for Novel Opportunities (GUANO)^{22–24} observations for the presence of high-frequency QPOs. Owing to the internal data readout time limitation of 0.1 ms, we could construct the light curve on this timescale, which provides a frequency coverage up to 5,000 Hz. The Leahy-normalized Swift/BAT PSD clearly shows broad signals with the peak power in the range 2,163–2,272 Hz (see Fig. 3a), and other signals with peak power in the range 1,464–1,614 Hz, consistent with all features we see in the ASIM-LED PSD. Swift/BAT GUANO observations, therefore, provide an independent confirmation of the $f_1 = 2,132$ Hz QPO seen with ASIM-LED.

Variations on timescales $t \leq 0.5$ ms can be caused by magnetic field instabilities in the magnetosphere close to the stellar surface (stellar radii, $r \leq 100$ km). Independently of whether we consider a crustquake

or magnetospheric instabilities as the initial trigger of GRB 200415A, Alfvén waves will be created in the magnetosphere. The Alfvén crossing time, $t_A \approx \pi r/c \approx 1$ ms (c , speed of light in a vacuum), of waves constrained to this region determines the timescale of reconnection and is the typical timescale during which instabilities develop^{25,26}. While bouncing back and forth between the footpoints of their respective magnetic field lines, these waves interact with each other nonlinearly across different field lines, dissipating energy. These ‘encounters’ produce emissions with a typical variability on a timescale of 1 ms, which is consistent with f_1 at 2,132 Hz. This process could be very inefficient and last for tens of Alfvén crossing times²⁷. $f_2 = 4,250$ Hz would then naturally be the overtone (within errors) of the broad feature at f_1 , or could be produced by the interaction of multiple Alfvén waves propagating along neighbouring field lines. After several crossings most of the energy of the waves has been absorbed by the crust. The reconnection is expected to terminate after a few milliseconds^{25,26}, which would naturally explain the sudden disappearance of the QPOs.

Furthermore, the observed timescale of $t \approx 1/f_1 = 0.469$ ms determines a lengthscale in which the energy of the giant flare is released, $l = c \times t \approx 140$ km. Assuming for simplicity a spherical magnetic field line, its radius would then be approximately 20 km. Hence the total volume of the flare is expected to be comparable in magnitude to, or somewhat larger than, the neutron star itself.

The properties of the observed spectra (Fig. 3, Table 1) are consistent with what is expected to happen during a reconnection event in the magnetosphere. The photospheric temperature of an expanding fireball, which forms as a consequence of the energy release during the reconnection, is expected to have typical temperatures in the range 20–80 keV (ref. 25), consistent with the 40–60 keV observed. The non-thermal component could be created by resonant cyclotron scattering, which dominates the scattering optical depth in highly magnetized magnetospheres and which can increase the energy of thermally emitted photons by orders of magnitudes. The strong magnetic field also enables us to understand the observed peak energy

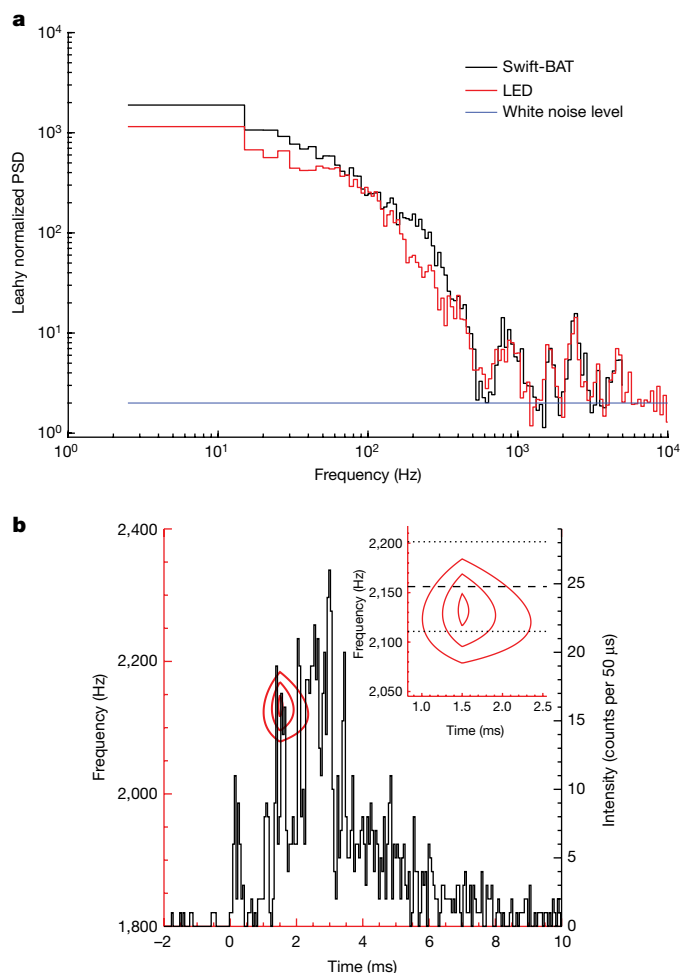


Fig. 3 | Periodogram and fits for quasi-periodicities search for the time interval 0–5 ms. **a**, Periodograms for ASIM-LED (50–400 keV; red) and Swift/BAT GUANO (15–150 keV; black) observations. The interval 0–200 ms is used for both periodograms. ASIM time resolution is 50 μ s (10-kHz upper frequency) and Swift/BAT time resolution is 100 μ s (5-kHz upper frequency). The blue horizontal line is the white noise level. The power spectral densities show broad signals after 700 Hz, providing independent confirmation that the $f_1 = 2,132$ Hz QPO determined from the ASIM-LED data is genuine. **b**, Dynamic power contours (in red) resulting from the Z^2 search in the 1,800–2,400 Hz frequency range, along with the LED light curve with 50- μ s time resolution. The inset is the expanded view of the Z^2 contours corresponding to the 99% (innermost), 95% (middle) and 90% (outermost) levels of peak Z^2 power centred at 2,132 Hz (Extended Data Table 3). For comparison the frequency found by the PSD analysis at $2,156 \pm 45$ Hz (the frequency and its uncertainty) is shown by dashed and dotted horizontal lines (Extended Data Table 2, time interval 0–10 ms).

around 1 MeV, because in very strong magnetic fields processes, such as one-photon pair creation and photon splitting, the energy of photons is limited to roughly $2m_e c^2$, where m_e is the mass of an electron.

An alternative explanation of the timing features is based on the proximity of the QPO candidate at $f_1 = 2,132$ Hz to one of the high-frequency QPOs observed in the tail of SGR 1806-20 with $f = 1,840$ Hz (ref. ⁹). The high-frequency QPOs in magnetars are commonly interpreted as radial overtones of the fundamental (magneto-)elastic oscillations with one or more nodes in the crust, which will be preferably excited during the flare²⁸. Depending on how exactly the instability in the magnetosphere is triggered, there may be strong perturbations in the crust of the neutron star which should naturally excite oscillations. Following this interpretation, the second strong feature at $f_2 \approx 4,250$ Hz may then be related to an even higher overtone. Within this interpretation

f_1 can be considered as an upper estimate on the purely shear mode with two (or three) nodes in the crust (see Methods section ‘QPO theoretical implications’). This interpretation is also consistent with theoretical expectations and the constraints²⁹ that were obtained from the QPOs of SGR 1806-20. However, the sudden disappearance of the QPOs after approximately 3.5 ms and the evolution of energy spectra slightly favour our first model, but do not exclude the presence of stellar oscillations.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-021-04101-1>.

- Kaspi, V. M. & Beloborodov, A. M. Magnetars. *Annu. Rev. Astron. Astrophys.* **55**, 261–301 (2017).
- Mereghetti, S., Pons, J. A. & Melatos, A. Magnetars: properties, origin and evolution. *Space Sci. Rev.* **191**, 315–338 (2015).
- Duncan, R. C. & Thompson, C. Formation of very strongly magnetized neutron stars: implications for gamma-ray bursts. *Astrophys. J.* **392**, L9–L13 (1992).
- Mazets, E. P., Golenetskii, S. V., Il’inskiy, V. N., Aptekar’, R. L. & Guryan, Yu. A. Observations of a flaring X-ray pulsar in Dorado. *Nature* **282**, 587–589 (1979).
- Ofek, E. O. et al. The short-hard GRB 051103: observations and implications for its nature. *Astrophys. J.* **652**, 507–511 (2006).
- Barat, C. et al. Fine time structure in the 1979 March 5 gamma ray burst. *Astron. Astrophys.* **126**, 400–402 (1983).
- Strohmayer, T. E. & Watts, A. L. Discovery of fast X-ray oscillations during the 1998 giant flare from SGR1900+14. *Astrophys. J. Lett.* **632**, L111–L114 (2005).
- Israel, G. L. et al. The discovery of rapid X-ray oscillations in the tail of the SGR1806-20 hyperflare. *Astrophys. J.* **628**, L53–L56 (2005).
- Strohmayer, T. E. & Watts, A. L. The 2004 hyper flare from SGR1806-20: further evidence for global torsional vibrations. *Astrophys. J.* **653**, 593–601 (2006).
- Pumpe, D., Gabler, M., Steininger, T. & Enßlin, T. A. Search for quasi-periodic signals in magnetar giant flares – Bayesian inspection of SGR1806-20 and SGR1900+14. *Astron. Astrophys.* **610**, A61 (2018).
- Marisaldi, M., Mezentsev, A., Østgaard, N., Reglero, V. & Neubert, T. GRB 200415A: ASIM observation. *GCN Circular 27622* (2020).
- Svinkin, D. et al. A bright γ -ray flare interpreted as a giant magnetar flare in NGC 253. *Nature* **589**, 211–213 (2021).
- Roberts, O. R. et al. Rapid variability of a giant flare from a magnetar in NGC 253. *Nature* **589**, 207–210 (2021).
- The Fermi-LAT Collaboration. High-energy emission from a magnetar giant flare in the Sculptor galaxy. *Nat. Astron.* **5**, 385–391 (2021).
- Yang, J. et al. GRB 200415A: a short gamma-ray burst from a magnetar giant flare? *Astrophys. J.* **899**, 106 (2020).
- Zhang, H. M., Liu, R.-Y., Zhong, S.-Q. & Wang, X.-Y. Magnetar giant flare origin for GRB 200415A inferred from a new scaling relation. *Astrophys. J. Lett.* **903**, 32 (2020).
- Burns, E. et al. Identification of a local sample of gamma-ray bursts consistent with a magnetar giant flare origin. *Astrophys. J. Lett.* **907**, 28 (2021).
- Neubert, T. et al. The ASIM mission on the International Space Station. *Space Sci. Rev.* **215**, 26 (2019).
- Østgaard, N. et al. The Modular X- and Gamma-Ray Sensor (MXGS) of the ASIM payload on the International Space Station. *Space Sci. Rev.* **215**, 23 (2019).
- Leahy, D. A. et al. On searches for pulsed emission with application to four globular cluster X-ray sources: NGC 1851, 6441, 6624 and 6712. *Astrophys. J.* **266**, L160–L170 (1983).
- Buccheri, R. et al. Search for pulsed γ -ray emission from radio pulsars in the COS-B data. *Astron. Astrophys.* **128**, 245–251 (1983).
- Gehrels, N. et al. The Swift gamma-ray burst mission. *AAstrophys. J.* **611**, 1005 (2004).
- Barthelmy, S. D. et al. The Burst Alert Telescope on the Swift Mission. *Space Sci. Rev.* **120**, 143–164 (2005).
- Tohuvavohu, A. et al. Gamma-Ray Urgent Archiver for Novel Opportunities (GUANO): Swift/BAT event data dumps on demand to enable sensitive subthreshold GRB searches. *AAstrophys. J.* **900**, 35 (2020).
- Mahlmann, J. F., Akgün, T., Pons, J. A., Aloy, M. A. & Cerdá-Durán, P. Instability of twisted magnetar magnetospheres. *Mon. Not. R. Astron. Soc.* **490**, 4858–4876 (2019).
- Parfrey, K., Beloborodov, A. M. & Hui, L. Twisting, reconnecting magnetospheres and magnetar spindown. *Astrophys. J. Lett.* **754**, 12 (2012).
- Li, X., Zrake, J. & Beloborodov, A. M. Dissipation of Alfvén waves in relativistic magnetospheres of magnetars. *Astrophys. J.* **881**, 13 (2019).
- Duncan, R. C. Global seismic oscillations in soft gamma repeaters. *Astrophys. J.* **498**, L45–L49 (1998).
- Gabler, M., Cerdá-Durán, P., Stergioulas, N., Font, J. A. & Müller, E. Constraining properties of high-density matter in neutron stars with magneto-elastic oscillations. *Mon. Not. R. Astron. Soc.* **476**, 4199–4212 (2018).

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2021

Methods

Instrumentation and data acquisition

Instrument description and observation geometry. The Modular X-ray and Gamma-ray Sensor (MXGS) is an imaging and spectroscopic X-ray and γ -ray instrument¹⁹ mounted on the starboard side of the Columbus module on the International Space Station. Together with the Modular Multi-Spectral Imaging Assembly (MMIA)³⁰ MXGS constitutes the instruments of the Atmosphere–Space Interactions Monitor (ASIM). The main objectives of MXGS are to image and measure the spectrum of X-rays and γ -rays from lightning discharges, known as terrestrial γ -ray flashes (TGFs). MMIA is designed to image and perform high-speed photometry of transient luminous events (TLEs) and lightning discharges. With these two instruments specifically designed to explore the relation between electrical discharges, TLEs and TGFs, ASIM is the first mission of its kind. Owing to telemetry limitations, ASIM is a triggered system; for each trigger, 2 s of data centred at the event is downloaded.

At the magnetar giant flare time, the ISS was 434.74 km above the Pacific, close to the New Zealand coast, at longitude 176.287° E and latitude 39.442° S. As the flare was localized⁶ by the Inter Planetary Network (IPN) on a region (a 4.73 arcmin \times 3.78 arcmin parallelogram centred at coordinates RA_{J2000}: 11.874°, dec_{J2000}: -25.194°) consistent with the NGC 253 galaxy (Extended Data Fig. 1), this was 71.25° off-axis, outside the MXGS imaging field of view (FOV). This is very close to the Earth limb, which covers 69.41° in the MXGS FOV, but it is high enough to rule out Earth's atmospheric absorption effects. Using the GEANT4 Monte Carlo simulation toolkit, considering the minimum altitude the photons reached (73.9 km) and the total column density traversed ($1.57 \times 10^{-3} \text{ g cm}^{-2}$) towards the direction of the NGC 253 galaxy, we see no effect on photons above 50 keV in terms of time delays or energy spectrum distortion, as the mean free path for 50-keV photons for this column density is $>10^6$ km.

Data processing and validation. From both HED and LED we have 2 s of data centred on the time of GRB200415A, with temporal resolution of 1 μ s (LED) and 28 ns (HED). For the analysis in this paper energies between 50 keV and 400 keV are used for LED (geometric area 1,024 cm²) and >300 keV for HED (geometric area 900 cm²). Although MXGS was designed for a different purpose, it provided unprecedented observations of the GRB 200415A magnetar flare, both regarding temporal resolution and spectral coverage. ASIM MXGS is the only instrument that recorded the main burst phase in the energy range from 50 keV up to 40 MeV without suffering from saturation effects (deadtime, pile-up). Here we describe the performance and limitations of the two instruments and emphasize that during the observations of the burst the observed flux and the measured energies were not affected by these limitations.

Instrument performance. It should be emphasized that both HED and LED are two independent instruments with independent read-out electronics. As for all instruments, HED and LED will miss counts when the flux exceeds a certain level. During GRB 200415A the flux never reached such a high level and here we explain why.

HED has 12 independent BGO/PMT bars. When an event occurs less than 10 μ s after the previous one in a single BGO/PMT bar, the second event will make a pulse that sits on the tail of the previous pulse, and is flagged as a 'fast' event. During the most intense part (0–3.6 ms) of GRB 200415A HED only recorded three 'fast' events out of a total of 393 raw counts.

LED has 16 independent readout chains and each chain has a dead-time of 1.4 μ s. When there is a hit in the same chain within 1.4 μ s the two hits will be recorded as a multi-hit event with one time tag. The pixel address will be lost and the two signals will be added, but the number of hits within the 1.4 μ s will be recorded. Hits in different chains will be recorded separately, but can have the same time tag if they occur within

the same 1 μ s. During the most intense part of GRB 200415A (0–3.6 ms) LED recorded 787 raw counts, of which 155 (19.7%) were multi-hits or had the same time tag. We use the interval 0–3.6 ms to ensure we include the spike at 3.4–3.6 ms (see Fig. 1b). If this fraction of 19.7% is due to Compton scattering in the detector, detector assembly or structures around the detector, they should only be counted as single hits for the purpose of timing analysis.

The ASIM mass model, developed using GEANT4, includes a detailed representation of the detectors as well as all surrounding structures and is well suited for testing how many Compton events should be expected. Incoming photons during the 0–3.6 ms interval with the spectral distribution given in Table 1 and shown in Fig. 3 were simulated and entered into the mass model from the correct incoming direction. It was found that 11–13% of these photons would deposit more than one count in the LED, and consequently should only be counted as single hits.

We also performed a probability analysis for the same interval (0–3.6 ms) to test how many real counts we could have missed owing to the 1.4- μ s deadtime in each single chain. The raw data were binned in 30- μ s bins, and the occurrence rate in each of the 120 bins (0–3.6 ms) was used to estimate how many real counts we will miss owing to the 1.4- μ s deadtime in each chain for each time bin. This analysis showed that we might have missed a total of about 14 counts in this interval, and they would have been recorded as multi-hits. If we consider that 14 of the multi-hits should have been counted as two real hits and Compton multi-hits reduced accordingly, there would be 17.6% Compton hits (instead of 19.7%) in this time interval, which is a slightly larger fraction than the expected 11–13% determined by using the mass model.

Accepted data. It is not possible to distinguish the possible real counts (among the multi-hits or with same time-tag in LED) from Compton scattered hits, and so we have excluded all the counts that are possible Compton photons from the analysis (except for spectroscopy, where raw counts are used). These Compton counts will also result in a non-Poissonian background distribution. To remove these counts, we proceeded as follows: Counts that are separated by less than 200 ns in HED are only counted as one. Counts that have the same time tag in LED (that is, within the same μ s) as well as multi-hits are also counted as only one count. This procedure will ensure that all Compton events will only be counted as single events, and also any cosmic showers or particle cascades in the 2 s of data will only be counted as single counts. After applying this procedure, the white noise (background) has a Poissonian distribution, and applying Leahy normalization²⁰ for the PSDs for both LED and HED indicates that the white noise has the expected constant power of 2. For the most intense part of the burst (0–3.6 ms) we end up with 363 and 663 accepted counts in HED and LED, respectively. For all plots in this paper, except for the spectroscopy analysis and the spectra shown in Fig. 3 and Extended Data Fig. 3, only accepted data are used.

ASIM was the only instrument that detected the peak intensity of the GRB in the energy range from 50 keV to tens of MeV: In Extended Data Fig. 2 we show GRB 200415A as measured by Fermi, ASIM and Swift/BAT. In the two upper panels it can be seen that the peak of GRB 200415A between 1.6 ms and 3 ms was clearly detected by both HED and LED, but not by the Fermi/GBM BGO or NaI. The lower panel shows that ASIM LED and Swift/BAT both detected the peak of GRB 200415A in energies below 350 keV, but ASIM is the only instrument that recorded the main burst in the high energy range (300 keV to 40 MeV). In Methods section 'QPO analysis', it will be shown that it is indeed in the rising phase of GRB 200415A that specific frequencies were observed by both LED and HED, with a statistically significant power level above the white noise level.

Time lag analysis. Following common practice for GRB studies we also tested whether there is an energy dependence in the arrival times of photons, especially during the initial hard spike emission of the giant flare. A cross-correlation analysis for the interval -2 ms to 10 ms using

50 μ s bins and steps showed that there is no time lag between the two light curves. Time lag (HED before LED) can only be seen if we narrow the bin size and steps down to 4 μ s or less.

Spectroscopy

Spectral analysis. To study the spectral evolution of the burst, we divided the time series into four intervals, as mentioned in the main text. For each time interval, we extracted the spectral files for LED and HED. For LED we used all counts, excluding multi-hits events, for which the energy information is accumulated over all simultaneous counts. For HED we used all counts except FAST events (see instrument performance) and all counts that are closer to the previous one in the same detector than a specific ‘safety time’. This time depends on the detector and the energy of the previous count, and ranges from 2 μ s to about 20 μ s for previous counts with energy above 20 MeV. This procedure has been introduced to ensure the highest reliability of the HED energy measurements, which are affected to some extent by the signal of the previous count. The procedure has been developed for the spectral analysis of terrestrial γ -ray flashes (TGFs), which exhibit fluxes much higher than those recorded for this event and therefore suffer much more from this effect. This procedure removes 53 out of a total of 2,141 counts above 400 keV (2.5%) and therefore does not substantially impact the spectral analysis. We do not use accepted data for the spectral analysis because the Compton events, counted as one by the accepted data procedure, are properly accounted for in the detector response matrices (DRM) and therefore are relevant for an accurate spectral analysis. We generated the DRM of the HED and LED detectors for the incoming directions of the burst. DRMs are built using the full mass model of the payload and the surrounding structures (the Columbus Module on the ISS). We then use XSPEC v12.10.1f, a widely used general-purpose X-ray spectral-fitting programme that is part of the XANADU package, for spectral fitting using the standard forward-folding approach. We fit all the time intervals in the energy range 60–350 keV for LED and 0.4–20 MeV for HED. Background is accumulated over the time interval -700 ms to -20 ms. Data were rebinned in order to have a minimum 1σ significance above the background in each bin, or a maximum of ten channels rebinned together. This latest condition is applied only in the highest energy part of the HED spectra. Rebinning is implemented on data files using the grppha utility from the FTOOLS package (https://heasarc.gsfc.nasa.gov/docs/software/ftools/ftools_menu.html). We test different models, including a power law with exponential cut-off (CPL), a simple power law (PO), and a blackbody spectrum (BB). The CPL is a custom model of the form $f(E) \approx E^\alpha \exp[-\frac{E(\alpha+2)}{E_p}]$ where α is the photon index and the peak energy E_p is the maximum of the spectral energy distribution. All the intervals can be fitted with a CPL. The best-fit parameters are shown in Table 1, and the fitted energy spectra and the corresponding residuals are shown in Extended Data Fig. 3. The precursor can also be well fitted with a blackbody spectrum with temperature 66^{+15}_{-12} keV, resulting in a reduced χ^2 comparable to that of the CPL model but only one free parameter; therefore, this model should be preferred. The preferred values for all the four phases are marked with bold in Table 1.

The spectral evolution can be summarized as follows: The spectrum of the precursor phase reveals a blackbody temperature of 66 keV. During the peak phase (0.8–3.2 ms) the flux increases by an order of magnitude and the spectrum is dominated by an accelerated component with a peak energy of 1,160 keV. It is during this acceleration phase that the two significant QPOs appear. The accelerated component dominates also in the decay phase, but the flux has decreased. In the tail phase the blackbody component reappears with a slightly lower temperature (40 keV) than in the precursor phase and the accelerated component is weaker.

Comparison with Fermi GBM. For comparison and validation purposes we compared the ASIM MXGS spectra with the Fermi GBM results

provided in ref. ¹³. We produced the LED and HED spectra for the same time intervals used in the Fermi analysis, and fitted the datasets with a cut-off power law model as done for Fermi. The time intervals are similar, but not identical to, those chosen in this paper. Extended Data Fig. 4 shows the ASIM unfolded data and the 1σ confidence interval model spectra for Fermi presented in Fig. 1d of ref. ¹³. Note that the unfolded data depend on the observed data, the response of the instrument and the considered theoretical best-fit model. Extended Data Table 1 reports the spectral parameters and the calculated fluxes for both ASIM MXGS and Fermi GBM. We note an overall good agreement between ASIM and Fermi spectral parameters. In particular, the peak energy is compatible, within errors, for both instruments for all time intervals, supporting the goodness of the calibration and spectral performance of ASIM MXGS. However, the calculated fluxes, although well compatible for intervals 1 and 4, are substantially different for intervals 2 and 3, ASIM reporting a much larger flux than Fermi. We note that these are the time intervals with the highest flux. In ref. ¹³ it is reported that interval 2 is affected by saturation and includes a correction factor for the flux obtained by comparing the GBM and Swift/BAT flux in the same time interval. We suggest that this correction may be underestimated, and saturation may be affecting GBM data in interval 3 as well. Conversely, we confirm that saturation in the GBM data does not significantly affect the overall spectral shape, as reported in ref. ¹³. We also note, in interval 4, a significant discrepancy between LED data and the GBM spectrum at energies below 300 keV. A simple cut-off power law fit does not yield a very good result, suggesting the need for an extra component at low energies that we model with a blackbody spectrum.

QPO analysis

Methodology. We approach the problem of detecting periodic or almost-periodic signals in noisy time series within a Bayesian framework^{31,32} in modelling the observed periodogram by assuming any low-frequency broadband variability to be due to a noise process. Owing to the lack of knowledge of the burst emission mechanism, we make a simple but conservative assumption that all broadband power in the periodogram is supplied by a red (aperiodic) noise process in the form of a power law. As opposed to QPO features—assumed to be fairly narrow features on top of this process—our assumption will cause weak signals at low frequencies to be buried in the higher variance of the broadband noise. In return, it would yield a very low false positive detection rate³³.

As in ref. ³¹, the periodogram refers to the squared Fourier transform of the data, and it is assumed to be a sample of the underlying physical process (the burst envelope and one or several noise processes). We make use of the Stingray Python Package³³ to perform this analysis.

We computed light curves and periodograms for LED and HED, for several time segments after the onset of the burst: 0–5 ms, 0–10 ms, 0–50 ms and 0–100 ms. A larger time interval gives much better frequency resolution. In each case, we produced a light curve by binning the data to a time resolution of $1/2\nu_{\text{Nyquist}} = 50 \mu\text{s}$, corresponding to a Nyquist frequency of $\nu_{\text{Nyquist}} \approx 10^4$ Hz. We chose this time resolution since, on the basis of typical neutron star sizes³⁴, we do not expect any QPO above 10 kHz.

We search both the unbinned periodogram as well as the same periodogram binned (to 3 and 5 times the original frequency resolution of 10^6 Hz). Additionally, we smooth the spectra with a Wiener filter with different smoothing factors (3 and 5) and compare results of the search of binned periodograms with searches across the smoothed periodograms.

An advantage of such a Bayesian approach is to provide a statistically rigorous framework to test whether additional model components (such as Lorentzian QPOs) are required by the data in the search for periodicities and quasi-periodicities. The Bayesian procedure has three main parts:

1. Broadband noise model selection. We address this task as a model-selection problem, within a fully Bayesian analysis, to find the

preferred broadband noise model to represent the periodogram's low-frequency part between two nested models. The first and simpler model, our null hypothesis H_0 , is just a power law to account for the red noise, plus a constant to account for the Poisson noise in the detection process. The second and more complex model, the alternative hypothesis H_1 , is a broken power law (see equation 8 in ref. ³²).

We used uniform prior probability distributions for all model parameters since no a priori knowledge of the magnetar giant flare emission is possible. We obtained the joint posterior distribution of given parameters using Bayes' theorem, taking the maximum a posteriori (MAP) estimates of the model parameters as our choice for the parameter estimation. We then performed a suit of Markov chain Monte Carlo simulations (MCMCs) within the Bayesian framework to ensure that the dataset is adequately described by the null hypothesis.

We compute the likelihood ratio test (LRT) statistic between the two models, or T_{LRT} , an often-used statistic for nested models, based on the ratio of the likelihood maxima values. We use 500 MCMC ensemble walkers with 100 samples each, after a burn-in phase with 200 samples for each walker. Then we can generate predictive (simulated) periodogram data from T_{LRT} , a 1,000 sub-sample of the parameter set, and compute a distribution for the statistic T_{LRT} to compare with observed $T_{\text{LRT}}^{\text{obs}}$ to obtain the P value. Given the distribution of T_{LRT} , we can compute the corresponding tail area probability, or the probability of obtaining a value of the test statistic as extreme as the one observed, under the assumption of the null hypothesis. The simulated data are fitted with both models H_0 and H_1 , following the same procedure performed on the observed data.

In all cases the preferred model is the simple power law, and we adopt model H_0 for the rest of our analysis of this burst.

2. Searching for (quasi-)periodicities. We search the periodogram for the highest data/model outlier and compare this outlier to those distributed by pure broadband noise to find narrow features that may be candidates for a possible QPO (see Extended Data Figs. 5b, 6b).

We use the chosen broadband noise model fit to the periodogram in the previous step to draw a new posterior predictive sample of the parameter sets, and simulate 1,000 periodograms. In the same fashion as outlined in the noise model selection, from these simulated periodograms (all following the chosen noise model), we compute the distribution of two other statistics, generated from the sample of simulations, to compare the observed value of the statistic and compute the corresponding P values.

One statistic is the summed-square residuals T_{SSE} , which is an indicator for how well the model fits the data. The second statistic is the maximum ratio of observed to model power, $T_{\text{R}} = \max_j \hat{R}_j$, where $\hat{R}_j = 2I_j/S_j$, and I_j and S_j are the observed and model powers. This is a sensible statistic with which to investigate the narrow features³⁵. The factor of two normalizes the residuals in such a way that \hat{R}_j will be distributed as χ^2 (see ref. ³⁵).

We search for the highest data/model outlier, T_{R} in the unbinned, binned and smoothed periodograms (see results in Extended Data Table 2). We can see how the features around 2,000–2,200 Hz and 4,000–4,400 Hz appear clearly and constantly. The few cases where the maximum ratio is detected at higher frequencies correspond with a higher frequency resolution, and with a smoothness factor of 3 or 5. So that re-binning the PSD to 3 or 5 times the original resolution is enough to detect these main features again. Because 4,256 Hz is consistently twice that of the former periodicity, within errors, it is probably the first harmonic of the broad feature at 2,156 Hz (both frequencies are marked with bold in Extended Data Table 2, LED 0–100 ms).

3. Search for QPOs in the binned data using the identical approach as for the model selection in the first step. We approach the search for QPOs in the data by following the same procedure applied to the broadband noise model selection. Thus from a model-selection point of view, and assuming a quasi-periodicity as another type of random process, we fit the periodogram simultaneously with a broadband noise process as well as one or more Lorentzians representing the QPO, as

the alternative hypothesis (H_1), against the broadband noise model only, which is the null hypothesis (H_0).

We use the residuals of the data divided by the preferred broadband noise model as a starting point from which to choose the alternative hypothesis. This is to decide whether to fit one or more Lorentzians together to model the broadband noise process.

Then we repeat the same procedure as above. We generate a large number of simulated periodograms from the MCMC sample of the posterior distribution of the parameter set, under the assumption of the null hypothesis (only noise). We compute then the LRT statistic from each simulated periodogram in turn, and finally compare the observed value of the statistic to the distribution generated from the sample of simulations, to get the corresponding P value associated with this model-selection problem; see Extended Data Table 2.

Many P values for the LRT test are much smaller than 10^{-6} . To limit the computational costs, they were estimated by the upper tail probability of the simulated posterior distributions, given the observed LRT, and corrected by the bin frequency size for each time interval and binning. Thus, the threshold for significance is adjusted for the number of different segments searched and the different binning.

Even so, the resulting P values (see Extended Data Table 2) are very small. The observed reduction in the likelihood ratio between null and the alternative hypothesis is larger than can be expected by chance if the null hypothesis (data were only broadband noise) were true, and we can reject it. However, this test cannot be seen as direct evidence that the alternative hypothesis is true³¹.

Moreover, to check the practical significance, we compute the effect size in terms of the discrepancy between the 95th percentiles of the simulated LRT distribution and the observed one (since only the upper tails of the distributions are of interest). This is $\varepsilon = |\text{LRT}^{\text{obs}} - \text{LRT}_{95}^{\text{sim}}|/\text{LRT}^{\text{obs}}$. This effect size was always larger than 0.85—that is, a relative distance in the percentile at least in a decrease of 15%, which agrees with the P value computation. The original data have a much higher discrepancy value or LRT statistic than any replicated datasets under the null hypothesis.

Another possible effect size could be comparing the fractional root-mean-square (RMS) amplitude in the QPO candidates, compared with the RMS in the simulated periodograms, under the broadband noise null hypothesis, over the same frequencies. This is $\varepsilon = |\text{RMS}^{\text{obs}} - \text{RMS}^{\text{sim}}|/\text{RMS}^{\text{obs}}$. The effect size is larger than 0.6 for the components around 4,250 and 2,132 Hz for LED data—that is, a relative distance in the simulated noised components RMS in a decrease of 40% (or an observed RMS 2.5 times the simulated). In HED, simulated RMS is over the 60% decrease (or an observed RMS above 66%). These effect sizes combined with the estimated P values support the existence and significance of these QPO candidates around the 2,132 and 4,250 Hz frequencies.

On the Z^2 search and significance of oscillations. We have also investigated the time evolution of the detected QPO features by employing the Z^2 statistics²¹. For this purpose, we used unbinned event arrival times and selected the number of harmonics as one—which is equivalent to the Rayleigh statistics—to search for a periodic signal³⁷. We first performed the search in the broad frequency range from 300 Hz to 5,000 Hz with the frequency resolution of 1 Hz using both LED and HED data collected within the first 5-ms time interval. We present the resulting Z^2 power spectra in Extended Data Fig. 7. Owing to the computational costs, we limited the finer resolution and dynamic search within 600-Hz frequency windows centred around the frequencies of detected QPOs from the PSD analysis. We performed the Z^2 search for each frequency with the data of LED and HED individually but within the same time intervals: we searched a time segment of 5 ms in the time interval ranging from $T_0 - 6$ ms to $T_0 + 50$ ms, by moving the search segment by 1-ms time steps.

Our Z^2 search with the LED data around the frequency 2,156 Hz resulted in the highest power at a frequency of 2,132 Hz. We calculated

Article

the significance of this signal as follows. Note that the powers of Z^2 with one degree of harmonic follows a χ^2 distribution with two degrees of freedom. We first calculated the random occurrence (chance) probability of the highest Z_1^2 power for a single trial case. We then computed the joint probability of having a Z^2 power spectrum of 1,000 frequency bins, which is the number of frequency steps used in our search and serves as the number of trial frequencies. We find that the 2,132 Hz signal is significant beyond the 99.9% confidence level both in LED and HED (Extended Data Fig. 7).

We calculated the chance occurrence probability for each quasi-periodic signal seen with the LED and HED data following the same method, and present them in Extended Data Table 3.

Duration of the QPOs. To investigate the duration of the high-frequency QPOs we inspected periodograms of the GRB region within ± 300 ms from the GRB onset. A time window of length 200 ms was used to produce Leahy-normalized PSD periodograms, which were combined into a spectrogram with time step of 1 ms, presented in Extended Data Fig. 8a. Time axis corresponds to the left edge of the 200-ms long sliding window. The sharp edge of the GRB spectral contents can be seen at around -200 ms, when the right edge of the sliding window just enters the GRB region. Another sharp edge at around 0 ms marks the moment when the left edge of the sliding window starts leaving the GRB region. Red noise is clearly seen up to 1 kHz. Before and after these sharp edges only white noise is seen, except for frequencies < 100 Hz in the tail of the GRB (lasting for ~ 100 ms).

On the magnified segments of the spectrogram focused on frequency range of 1 kHz to 5 kHz, and produced with a higher time resolution (10- μ s time step), the appearance and disappearance of the high-frequency QPOs can be seen (Extended Data Fig. 8b, c). A clear transition region with strong perturbations is seen between 1 ms and 3.5 ms, where the main fine time structure of the burst phase is concentrated. The motion of the sliding window across this area shows the spectral evolution of the signal contents. After ~ 3.5 ms the 2,132-Hz QPO smears out down to the white noise level, and we will refer to 3.5 ms as the end of the QPO.

QPO theoretical implications

As discussed above, our analysis indicates two statistically significant QPOs, one at $f_1 = 2,132$ Hz and the other at $f_2 = 4,250$ Hz; both seen only during the peak phase (0–3.2 ms). Therefore, in this section, we focus on the theoretical interpretation of these two frequencies. There are two distinct physical processes that may result in signals with frequencies in this range: oscillations of the neutron star excited during the giant flare and magnetic reconnection events in the magnetosphere.

Neutron star oscillations. There are a number of different oscillation modes in neutron stars that may give rise to frequencies above 1 kHz (ref. 34): pressure modes, fundamental modes, space-time modes, superfluid modes and magneto-elastic modes. Of these, the latter have the strongest potential to create an observable temporal variation of the electromagnetic emissions during a magnetar giant flare^{28,38–40}. Additionally, they are expected to be preferentially excited during the flare²⁸. The strongest support in favour of this explanation is the detection of this kind of QPO in the decaying tail of previously observed giant flares^{7–9,41–43} of SGR 1806-20 and SGR 1900+14, although some doubt was cast on the significance of these detections recently¹⁰. The QPOs can be divided into two groups with low ($f \lesssim 150$ Hz) and high ($f \gtrsim 500$ Hz) frequencies. In particular, the 1,840-Hz oscillation detected⁹ in SGR 1806-20 is very close to f_1 . The low-frequency QPOs are usually related to magneto-elastic oscillations without nodes ($n = 0$) in the radial direction inside the crust, whereas the high-frequency QPOs could be oscillations that have one or more nodes ($n \geq 1$) inside the crust.

To estimate the properties of the neutron star we have used the empirical formula for the different oscillation frequencies given in

ref. 29. If the QPOs indeed represent stellar oscillations, the high frequencies indicate that they are connected to the overtones of the magneto-elastic oscillations that have one or more modes in the crust, $l t_n$, where l indicates the angular dependency and n is the number of nodes inside the crust in the radial direction. The dependence of these oscillations on l is much weaker than the dependence on n , and we therefore concentrate on the latter and refer to the $l = 2$ QPO from now on. Even for the most extreme (most massive) neutron star models, none of the equations of state considered in the literature provides crustal shear mode frequencies for 2t_1 of the order of 2 kHz or more^{36,44}.

The frequencies of the $n = 1$ modes are predicted⁴⁴ to be approximately $f \approx 400$ –800 Hz, with very extreme cases from 200 up to 1,500 Hz. Therefore, it is reasonable to associate f_1 with a higher $n > 1$ overtone. The presence of a signature at $\sim 1,400$ Hz indicates that this may be associated with lower $n > 0$ modes. However, owing to its low significance we will not consider it further and we will discuss this elsewhere in more detail.

The second strong feature at f_2 may then be related to a higher overtone. If the interpretation in terms of neutron star oscillations is correct, this would limit the available choices of the crust model, because the order of the different overtones of the crustal shear modes is not simply 1:2:3:4:5, and thus to match f_1 and f_2 requires some fine-tuning.

However, in reality the magnetic field complicates this analysis and can increase the frequencies substantially²⁹. Note that the maximal frequency increase allowed when including magnetic fields limits the frequency of 2t_1 to about twice the frequency of the pure shear mode. Otherwise the magnetic field would dominate and no high-frequency QPO would be expected²⁹. Thus the magnetic field helps to increase the frequency, but cannot bring typical shear mode frequencies of 2t_1 up to more than 2 kHz. Therefore, f_1 can be considered as an upper estimate on the purely shear mode frequency of the 2t_2 QPO (in the case that the equation of state predicts that $f_1 > 2f_{2,t_2}$, the conclusion would be that f_1 is an upper estimate on f_{2,t_3}). Without independent estimates for the magnetic field or the properties of the star, the parameter space spanned by the equation of state, mass of the neutron star model, and magnetic field strength and configuration is too big to further constrain the properties of the neutron star.

Reconnection in the magnetosphere. Although the frequency $f_1 = 2,132$ Hz is intriguingly similar to the highest frequencies of the QPOs observed in SGR 1806-20, there is a fundamental difference in both observations. For SGR 1806-20 in particular the 150-Hz and 625-Hz QPOs were observed for several tens of seconds during the X-ray tail⁹—however, the data are consistent with short-lived signals < 0.5 s appearing and disappearing⁴³—whereas here the QPO candidate is only observed for ~ 3 ms during the initial short peak. The timescale $t \leq 1$ ms associated with $f_1 = 2,132$ Hz corresponds very well to the Alfvén crossing time t_A of the nearby magnetosphere, which is essentially equal to the light crossing time there. For the Alfvén waves confined to a region close to the stellar surface ($r \leq 100$ km) this means $t_A \approx 1$ ms. The Alfvén crossing time determines the timescale of reconnection and is the typical timescale during which magnetospheric instabilities develop^{25,26,45,46}. Therefore, it is tempting to assume that the fast variability of the emission on a timescale of 1 ms can be explained by Alfvén waves dissipating energy. Recently it has been suggested that dissipation in the magnetosphere owing to nonlinear interactions of Alfvén waves could be an inefficient process and could last for tens of Alfvén crossing times^{27,46}. Most of the energy in these waves is absorbed by the crust and re-emitted on longer timescales (weeks to months). In this picture, the initial trigger of GRB 200415A (which could be either a crustquake or a magnetospheric instability) generates Alfvén waves in the magnetosphere, which propagate along the magnetic field lines. These Alfvén waves bounce back and forth between the footpoints of the field lines, naturally defining the Alfvén crossing time as the typical timescale. When Alfvén waves travelling along different field lines

interact nonlinearly, reconnection of the field lines occurs and energy is released. These encounters produce emission with a typical variability on a timescale of 1 ms. The frequency $f_2 = 4,250$ Hz would then naturally be the overtone, or could be produced by the interaction of multiple Alfvén waves propagating along the neighbouring field lines. After approximately ten crossings most of the energy in the waves has been absorbed by the crust. The longer-lasting emission probably has a different origin, related to the evaporation of plasma trapped in the magnetosphere by the ultra-strong magnetic field⁴⁷.

Data availability

The ASIM data used for this study are available on the Zenodo repository, <https://doi.org/10.5281/zenodo.5242975>. The Swift 100- μ s-resolution data are publicly available at the Swift BAT GUANO repository: <https://www.swift.psu.edu/guano>. The Fermi NaI and BGO data are publicly available at <https://heasarc.gsfc.nasa.gov/W3Browse/fermi/fermig-brst.html>.

Code availability

For the search for QPOs within the Bayesian framework, we use the Stingray Python Package developed by Huppenkothen et al.³⁰ (<https://docs.stingray.science>). The Z^2 search and related calculations of confidence level and chance probabilities are performed using Interactive Data Language (IDL) version 8.7.3. The IDL codes developed can be provided upon request. For spectral analysis we used the software XSPEC, version v12.10.1f, (including the XANADU package) available at <https://heasarc.gsfc.nasa.gov/xanadu/xspec/>.

- Chanrion, O. et al. The Modular Multispectral Imaging Array (MMIA) of the ASIM payload on the International Space Station. *Space Sci. Rev.* **215**, 28 (2019).
- Vaughan, S. A Bayesian test for periodic signals in red noise. *Mon. Not. R. Astron. Soc.* **402**, 307–320 (2010).
- Huppenkothen, D. et al. Quasi-periodic oscillations and broadband variability in short magnetar bursts. *Astrophys. J.* **768**, 87 (2013).
- Huppenkothen, D. et al. Stingray: a modern Python library for spectral timing. *Astrophys. J.* **881**, 39 (2019).
- McDermott, P. N., van Horn, H. M. & Hansen, C. J. Nonradial oscillations of neutron stars. *Astrophys. J.* **325**, 725–748 (1988).
- Vaughan, S. A simple test for periodic signals in red noise. *Astron. Astrophys.* **431**, 391–403 (2005).
- Sotani, H., Kokkotas, K. D. & Stergioulas, N. Torsional oscillations of relativistic stars with dipole magnetic fields. *Mon. Not. R. Astron. Soc.* **375**, 261–277 (2007).
- Brazier, K. T. S. Confidence intervals from the Rayleigh test. *Mon. Not. R. Astron. Soc.* **268**, 709–712 (1994).
- Timokhin, A. N., Eichler, D. & Lubarsky, Yu. On the nature of quasi-periodic oscillations in the tail of soft gamma repeater giant flares. *Astrophys. J.* **680**, 1398–1404 (2008).
- Gabler, M., Cerdá-Durán, P., Stergioulas, N., Font, J. A. & Müller, E. Modulating the magnetosphere of magnetars by internal magneto-elastic oscillations. *Mon. Not. R. Astron. Soc.* **443**, 1416–1424 (2014).
- Gabler, M. et al. Magneto-elastic oscillations modulating the emission of magnetars. *Astron. Nachr.* **338**, 1105–1108 (2017).
- Watts, A. L. & Strohmayer, T. E. Detection with RHESSI of high-frequency X-ray oscillations in the tail of the 2004 hyperflare from SGR1806-20. *Astrophys. J.* **637**, L117–L120 (2006).

- Huppenkothen, D., Heil, L. M., Watts, A. L. & Göğüş, E. Quasi-periodic oscillations in short recurring bursts of magnetars SGR1806-20 and SGR1900+14 observed with RXTE. *Astrophys. J.* **795**, 114 (2014).
- Miller, M. C., Chirenti, C. & Strohmayer, T. E. On the persistence of QPOs during the SGR1806-20 giant flare. *Astrophys. J.* **871**, 95 (2019).
- Steiner, A. W. & Watts, A. L. Constraints on neutron star crusts from oscillations in giant flares. *Phys. Rev. Lett.* **103**, 181101 (2009).
- Carrasco, F., Viganò, D., Palenzuela, C. & Pons, J. A. Triggering magnetar outbursts in 3D force-free simulations. *Mon. Not. R. Astron. Soc.* **484**, L124–L129 (2019).
- Li, X. & Beloborodov, A. M. Plastic damping of Alfvén waves in magnetar flares and delayed afterglow emission. *Astrophys. J.* **815**, 25 (2015).
- Thompson, C. & Duncan, R. C. The giant flare of 1998 August 27 from SGR1900+14. II. Radiative mechanism and physical constraints on the source. *Astrophys. J.* **561**, 980–1005 (2001).

Acknowledgements This work made use of data from ASIM. ASIM is an international experiment aboard the International Space Station. ASIM is a mission of the European Space Agency (ESA), funded by ESA and national grants from Denmark, Norway and Spain. ESA PRODEX contracts C4000115884 (DTU) and 4000123438 (Bergen) support the ASIM Science Data Centre. The science analysis is supported by ESA Topical Team contract 4200019920/06/NL/VJ; European Research Council grant AdG-FP7/2007-2013: n 320839; Research Council of Norway contract 223252/F50 (CoE/BCSS); and Ministerio de Ciencia, Innovación y Universidades grants ESP 2017-86263-C4 and ESP2017-87676-C5-5-R. J.A.F. and P.C.-D. acknowledge support by the Spanish Agencia Estatal de Investigación (grant PG2018-095984-B-I00) and by the Generalitat Valenciana (PROMETEO/2019/071). P.C.-D. acknowledges the Ramon y Cajal funding (RYC-2015-19074) for supporting his research. M.G. acknowledges support through the Generalitat Valenciana via the grant CIDEAGENT/2019/031. A.J.C.-T., J.P.-G., M.D.C.-G. and E.R. acknowledge financial support from the State Agency for Research of the Spanish MCIU through the “Center of Excellence Severo Ochoa” award to the Instituto de Astrofísica de Andalucía (SEV-2017-0709). This work is partly based on observations made with the Gran Telescopio Canarias (GTC), installed at the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofísica de Canarias, on the island of La Palma. We thank J.D. Scargle for his help with the preliminary time-series analysis and his comments. J.P.-G. also acknowledges funding support from Spanish public funds for research from project PID2019-107061GB-C63 from the ‘Programas Estatales de Generación de Conocimiento y Fortalecimiento Científico y Tecnológico del Sistema de I+D+i y de I+D+i Orientada a los Retos de la Sociedad’ from the Spanish Ministry of Science, Innovation and Universities (MCIU).

Author contributions A.J.C.-T., N.Ø. and V.R. led the research team for this event. V.R., J.N.-G., N.Ø. and T.N. contributed to ASIM and provided the raw ASIM data. A.M., M.M., N.Ø., A.L. and V.R. provided the data for the time analysis. E.G., C.S.-G., J.P.-G. and E.R. analysed the time variability in the data. A.L. and C.A.O. provided the spectral accepted data. B.-B.Z. and Z.-K.P. performed the preliminary spectral analysis and M.M. performed the detailed spectral analysis. I.K. and C.B.-J. contributed to all ASIM phases for the MXGS development. A.L. and D.S. verified the LED and HED energy-response matrices. D.S. provided the calibration of the HED instrument and also performed the calculations of column density and potential photon absorption in the atmosphere. F.C., C.J.E., G.G., K.U. and S.Y. contributed substantially to the design, building and operation of the ASIM MXGS instrument. I.M.C.-G., M.A.C.T., E.F.-G., C.P.d.P., S. G. and A.J.R.T. provided the optical data. P.C.-D., M.G. and J.A.F. contributed to the theoretical interpretation of the observational results. A.J.C.-T., N.Ø., E.G., C.S.-G., M.G. and J.P.-G. contributed to the writing of the manuscript. All the co-authors, including S.B., M.D.C.-G., A.C., M.A.C.T., P.C.-D., J.A.F., M.G., Y.-D.H., I.K., N.L., A.N.G., S.B.P., R.S.-R. and T.S., have read the manuscript and made useful suggestions.



Competing interests The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-021-04101-1>.

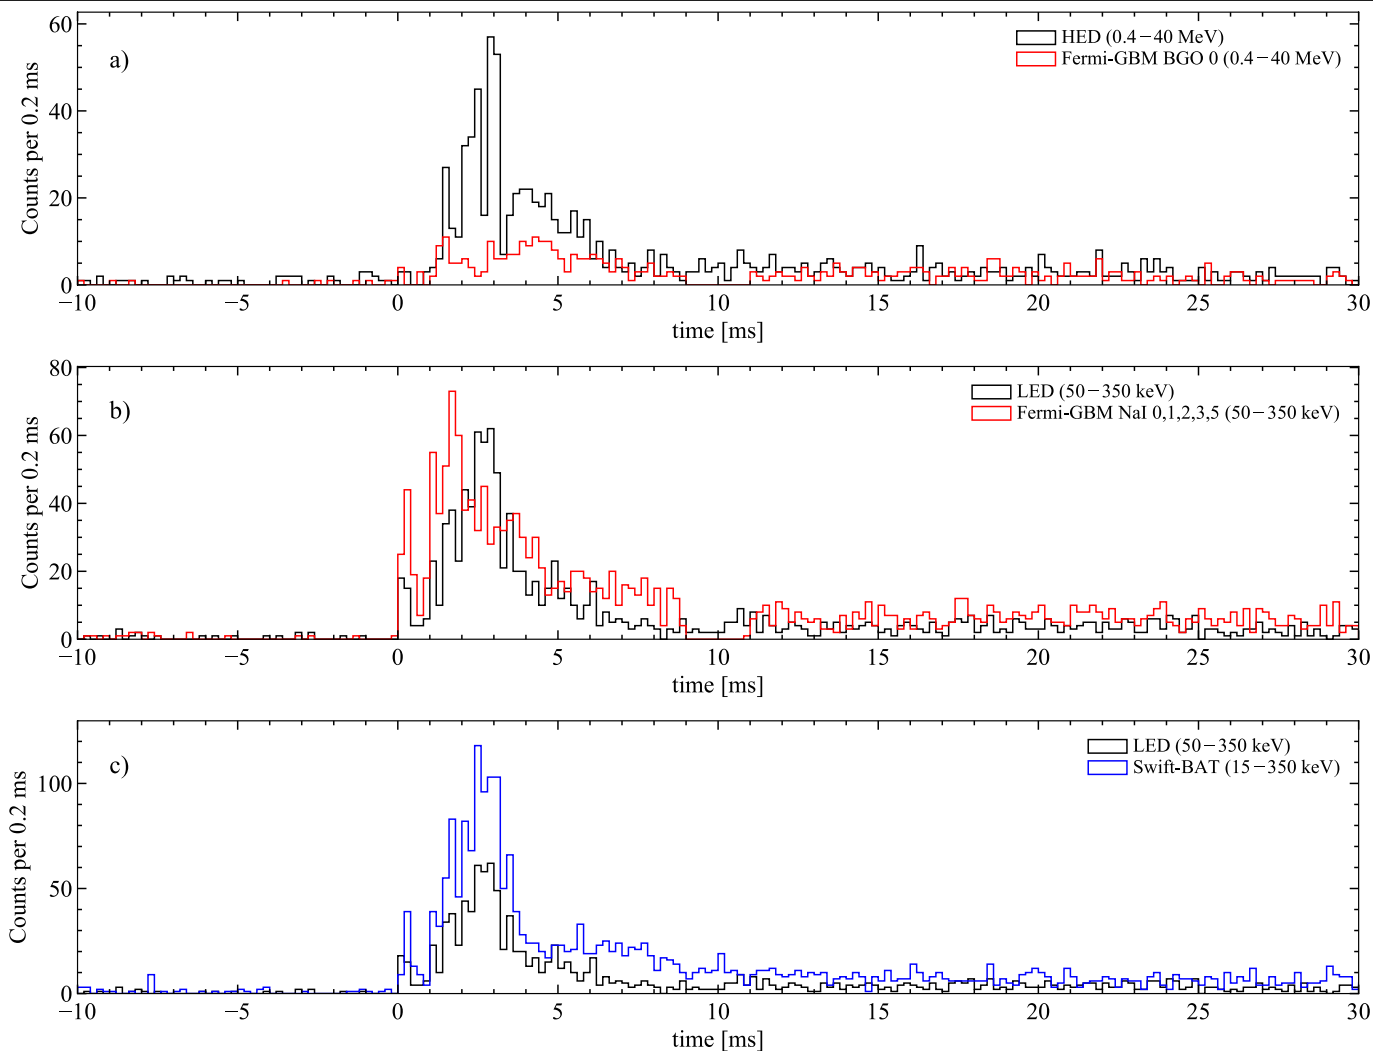
Correspondence and requests for materials should be addressed to N. Østgaard, E. Göğüş, A. Mezentsev, M. Gabler or M. Marisaldi.

Reprints and permissions information is available at <http://www.nature.com/reprints>.




Extended Data Fig. 1 | Optical content of the final Inter-Planetary Network (IPN) error box¹² (red parallelogram) partly including the NGC 253 galaxy. **a.** A 120-s exposure (clear filter) taken by the BOOTES-1A wide field camera in Huelva (South Spain) on 23 July 2020 (04:00 UT). The field of view is $0.5^\circ \times 0.5^\circ$.

This field was just 1.84° above the Earth's limb as seen by ASIM. **b.** A magnification for the IPN error box showing a colour image resulting from the combination of g (10 s), r (5 s) and i (5 s) -band exposures taken by the 10.4-m Gran Telescopio Canarias on 16 August 2020, 04:30 UT.



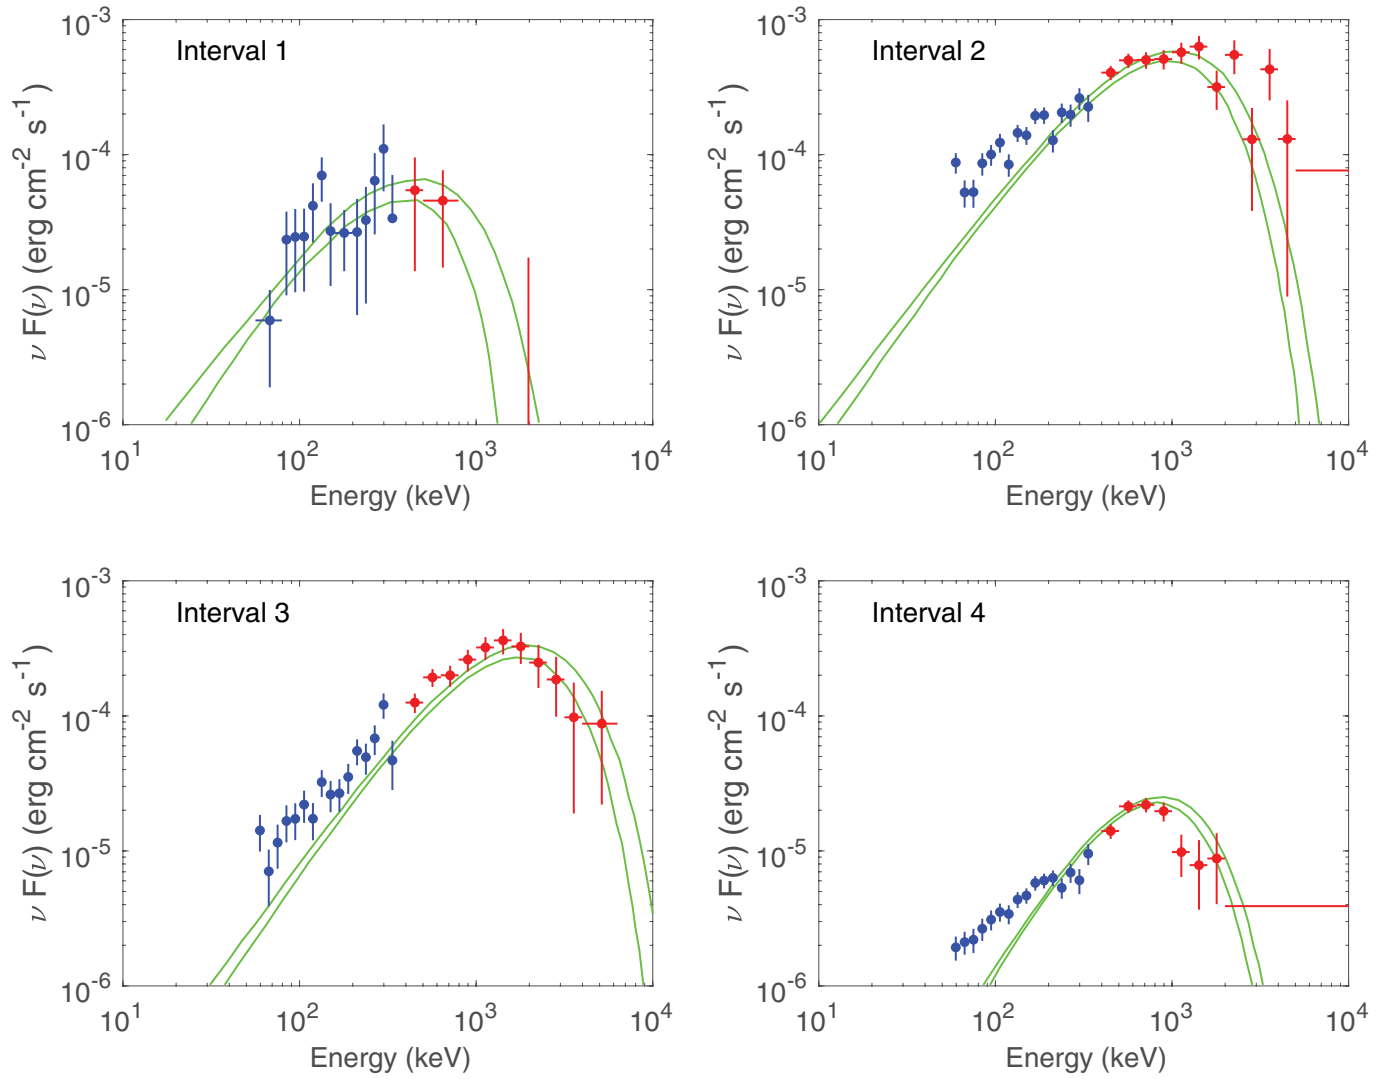
Extended Data Fig. 2 | GRB 200415A lightcurve as measured by Fermi/GBM, ASIM and Swift/BAT. a, ASIM HED (all 12 BGOs) and Fermi/GBM BGO 0 (as shown in ref. ¹³) in the same energy ranges. **b,** ASIM LED and Fermi/GBM NaI

0,1,2,3,5 (as in ref. ¹³) in the same energy range. **c,** ASIM LED and Swift/BAT in comparable energy ranges.



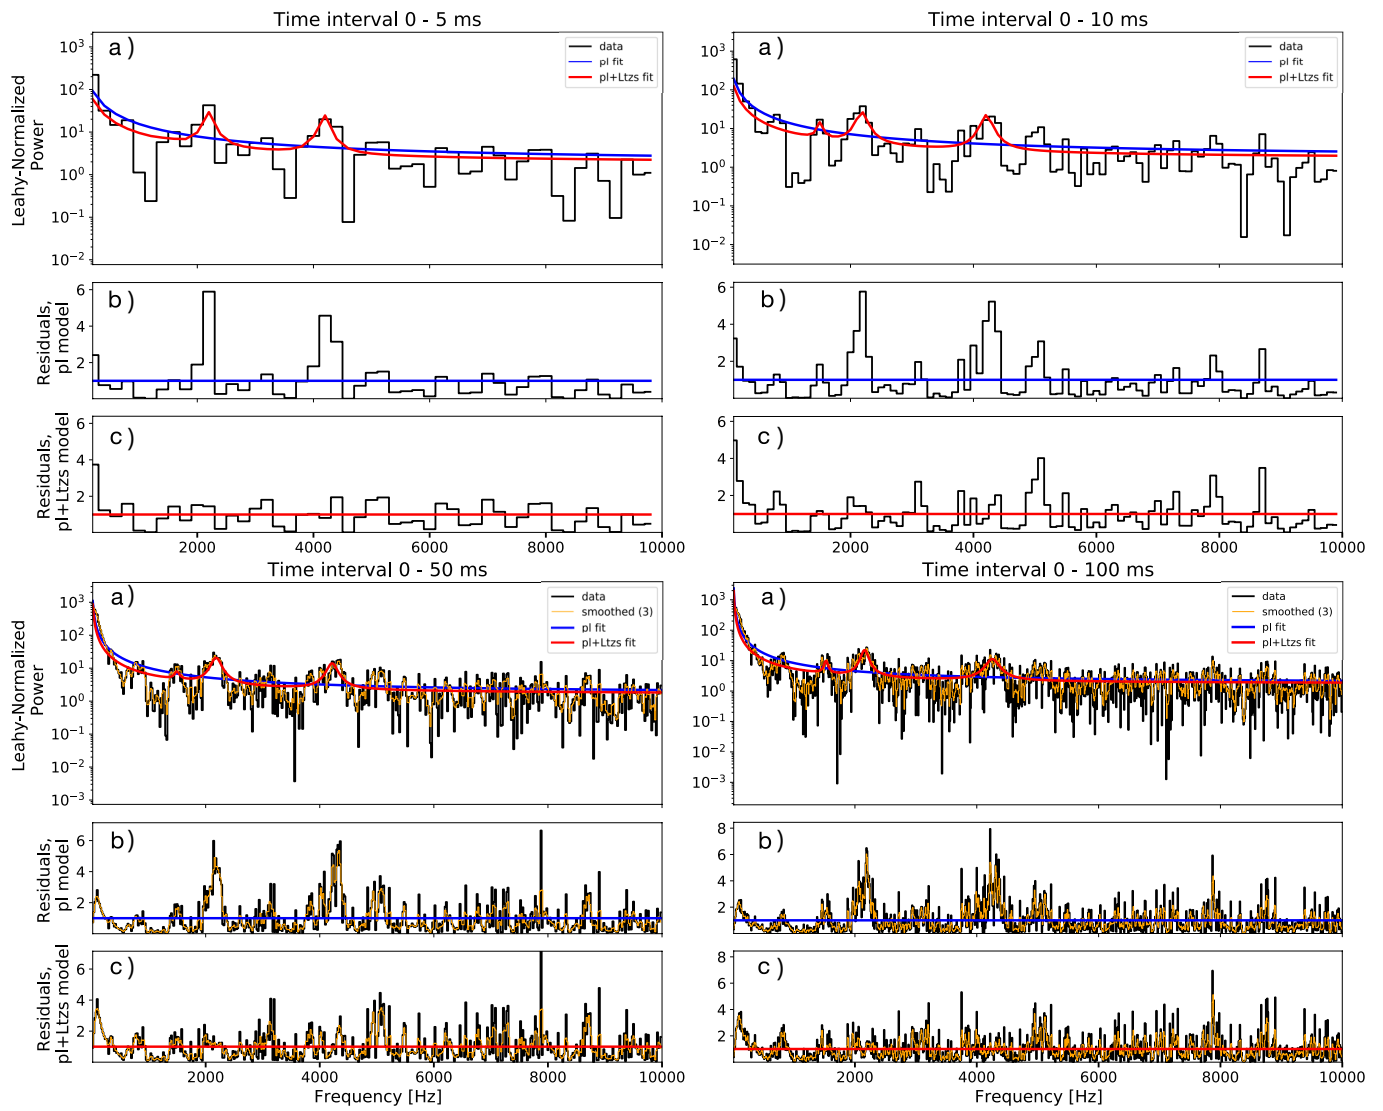
Extended Data Fig. 3 | Time-resolved spectroscopy (including residuals). Spectra for four time intervals during the magnetar giant flare, defined in the main text. **a**, Precursor: 0.0–0.8 ms. **b**, Peak: 0.8–3.2 ms. **c**, Decay: 3.2–8 ms. **d**, Tail: 8–160 ms. A total number of 2,148 counts in the LED and 2,141 counts in

the HED have been used. Left panels: count spectra (blue: LED; red: HED) and best-fit models (green). Middle panels: residuals. Best-fit models are those highlighted in bold in Table 1. Right panels: energy flux density for the best-fit models.



Extended Data Fig. 4 | Comparison between ASIM MXGS and Fermi GBM. Spectral energy density plots for the four time intervals considered in ref. ¹³. Data points: ASIM MXGS unfolded spectra; LED (blue) and HED (red). Green curves: 1σ confidence region for the Fermi GBM spectra, from figure 1d from ref. ¹³.

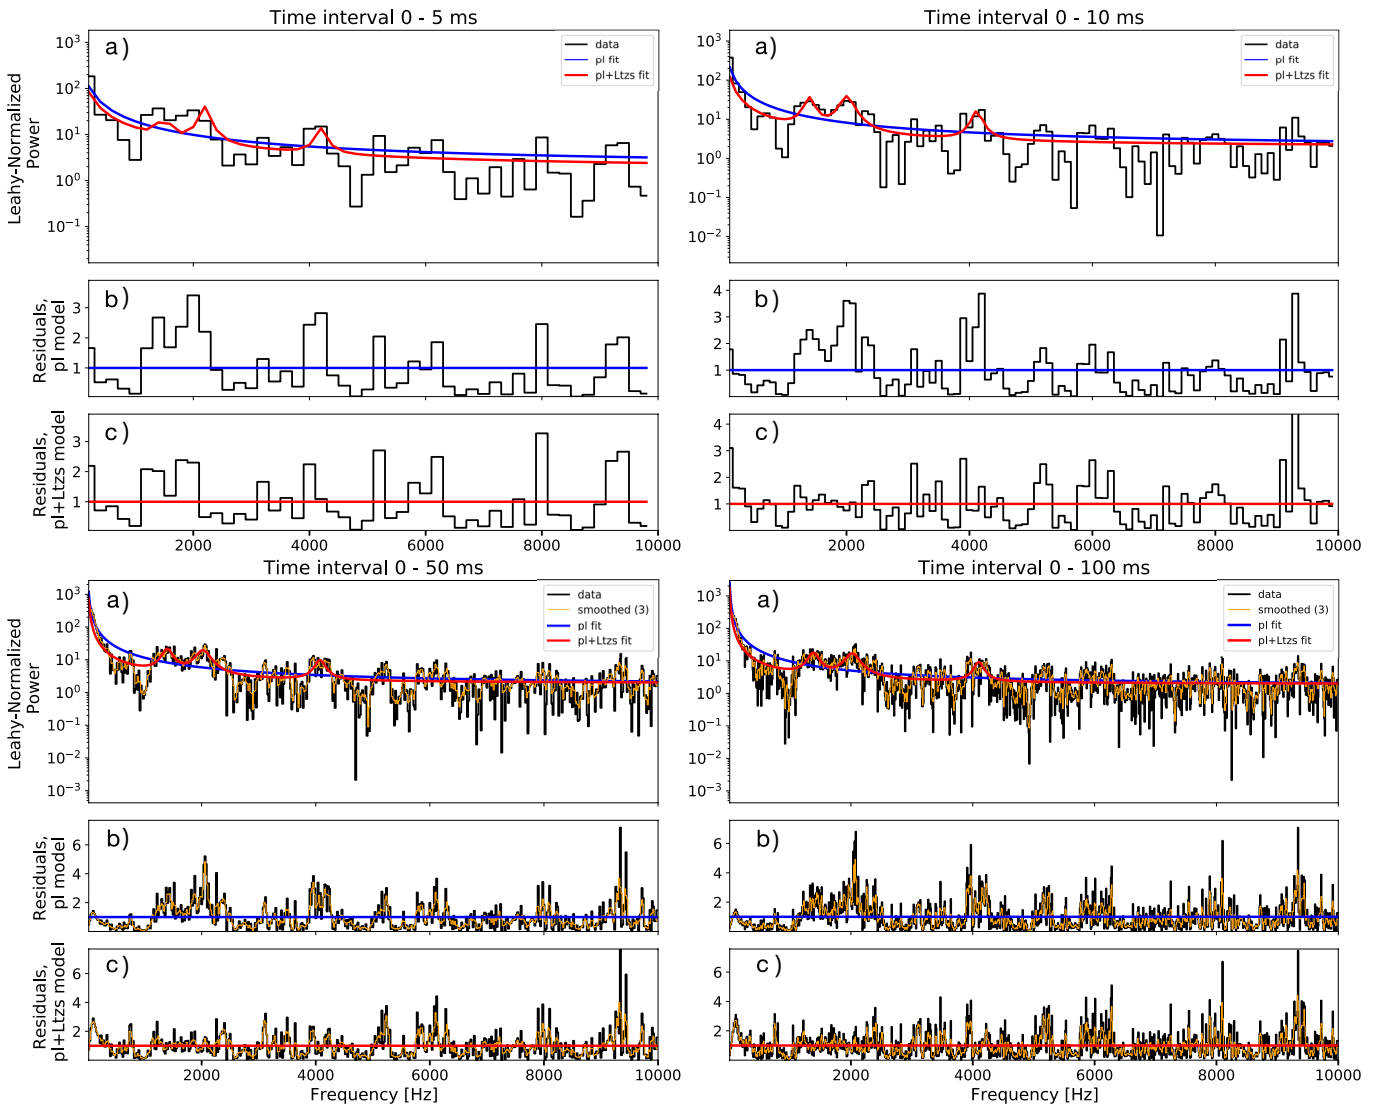
QPO Search (LED)



Extended Data Fig. 5 | Periodograms and residuals of the LED light curves for time periods from 0-5, 0-10, 0-50 and 0-100 ms after the burst ($t=0$ is start of the burst), with a time resolution of $50 \mu\text{s}$ (frequency bins of 200 Hz). a, Unsmoothed (black) and smoothed (orange; Wiener filter, $3\Delta\nu$) periodograms.


Power-law fit in blue, and power law plus two Lorentzians (centred around the 2,100 and 4,200 Hz features) fit in red. b, The residuals of the fits around the 2,100 Hz feature. c, The residuals of the fits around the 4,200 Hz feature.

QPO Search (HED)



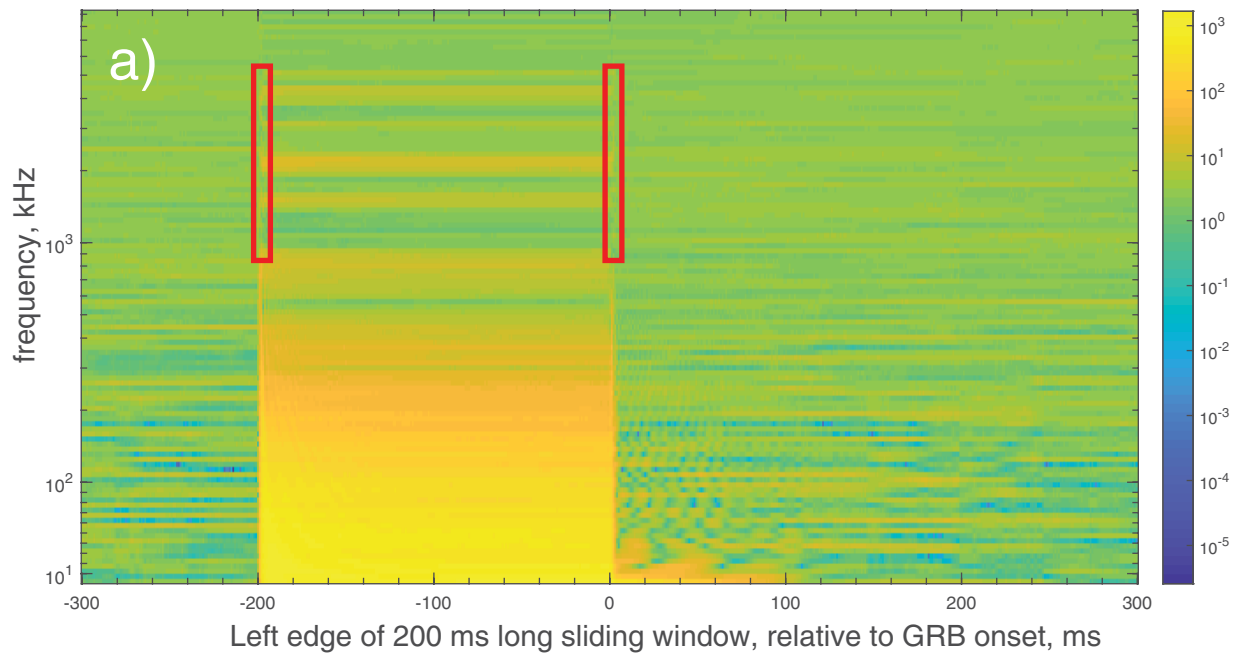


Extended Data Fig. 6 | Periodograms and residuals of the HED light curves for time periods from 0-5, 0-10, 0-50 and 0-100 ms after the burst ($t=0$ is start of the burst), with a time resolution of 50 μ s (frequency bins of 200 Hz). a, Unsmoothed (black) and smoothed (orange; Wiener filter, $3\Delta\nu$)

periodograms. Power-law fit in blue, and power law plus two Lorentzians (centred around the 2,100 and 4,200 Hz features) fit in red. b, The residuals of the fits around the 2,100 Hz feature. c, The residuals of the fits around the 4,200 Hz feature.



Extended Data Fig. 7 | Z_1^2 power spectra in the 300-5,000 Hz frequency range, obtained with the events within the first 5 ms time interval of the burst. a, Using LED data. b, Using HED data. The horizontal dotted, dashed and solid lines represent 95%, 99% and 99.9% confidence levels, respectively.



Extended Data Fig. 8 | A spectrogram combined of consecutive periodograms for ASIM-LED, 50-400 keV, with 200 ms long sliding window, and 1 ms step. Time axis corresponds to the left edge of the sliding window. **a**, Covering the interval between -300 ms and +300 ms relative to the GRB onset. Colour bar represents the Leaky-normalized PSD, where 2 corresponds to the white noise level. Areas of interest are highlighted with the red rectangles and are shown in the next panels for frequency range 1 kHz to

5 kHz. **b**, Zoomed-in segment of the spectrogram when the sliding window just enters the GRB region. A higher time resolution of 10 μ s steps is used. For convenience time axis corresponds to the right edge of the 200-ms long sliding window, relative to the GRB onset. **c**, Zoomed-in segment of the spectrogram when the sliding window starts to leave the GRB region. Time axis corresponds to the left edge of the 200-ms long sliding window, relative to the GRB onset.

Extended Data Table 1 | Spectral parameters and fluxes for ASIM and Fermi

Interval	Time ^a (ms)	ASIM - MXGS				Fermi - GBM		
		Red. χ^2 (d.o.f.)	Photon Index	Peak energy (keV)	Flux ^c 10^{-5} erg/cm ² s	Photon Index	Peak energy (keV)	Flux ^d 10^{-5} erg/cm ² s
1	-0.004 - 0.996	0.83 (13) ^b	$0.02^{+1.1}_{-0.83}$	337^{+96}_{-67}	$10.3^{+9.5}_{-4.9}$	-0.08 ± 0.23	428 ± 71	9.9 ± 1.2
2	0.996 - 3.596	1.6 (25)	$-0.86^{+0.11}_{-0.11}$	1182^{+105}_{-93}	$90.8^{+6.4}_{-6.0}$	-0.21 ± 0.08	997 ± 77	33.7 ± 1.5
3	3.596 - 7.396	0.82 (24)	$-0.36^{+0.18}_{-0.16}$	1590^{+160}_{-130}	49.7^{+15}_{-11}	-0.11 ± 0.08	1856 ± 155	17.5 ± 0.81
4	7.396 - 140.796	1.9 (21)	$-0.78^{+0.15}_{-0.14}$	1000^{+130}_{-100}	$2.65^{+0.29}_{-0.27}$	0.34 ± 0.08	846 ± 39	2.69 ± 0.06

^aTime referred to ASIM T_0 , corresponding to Fermi time intervals (ref. ¹³).

^bC statistics is used instead of χ^2 owing to the low number of counts available.

^cIn the 60 keV–10 MeV range, 68% confidence.

^dIn the 8 keV–10 MeV range, 68% confidence. Data presented in table 1 of ref. ¹³.

Extended Data Table 2 | Results for the search for periodicities and quasi-periodicities: posterior summary

			LED						HED					
			Periodicity Search			QPO Search			Periodicity Search			QPO Search		
time interval	bin size		freq	T_R	$p(T_R)$	freq	RMS	$p(LRT)$	freq	T_R	$p(T_R)$	freq	RMS	$p(LRT)$
(ms)	(Hz)		(Hz)			(Hz)	%		(Hz)			(Hz)	%	
0-5	200	unbin	2200	11.79	0.095 ± 0.021	1504.39 \pm 15.31	13.39 \pm 2.61	0.0032 ± 10^{-4}	2000	6.81	0.865 ± 0.024	1439.71 \pm 35.81	39.31 \pm 0.95	$6e-4 \pm 6 \times 10^{-5}$
						2193.91 \pm 50.88	25.44 \pm 2.54					2099.58 \pm 52.21	33.13 \pm 2.12	
						4200.76 \pm 28.02	20.79 \pm 1.77					4133.48 \pm 37.46	23.03 \pm 3.63	
0-10	100	unbin	2200	11.58	0.255 ± 0.031	1501.48 \pm 6.69	15.41 \pm 3.24	0.0016 ± 10^{-4}	4200	7.75	0.915 ± 0.019	1442.59 \pm 187.14	41.96 \pm 0.89	0.005 ± 0.005
						2156.11 \pm 45.33	30.00 \pm 2.16					2087.77 \pm 406.43	39.32 \pm 1.79	
						4255.49 \pm 211.86	24.09 \pm 1.53					4128.98 \pm 217.98	25.56 \pm 3.27	
0-50	20	unbin	7880	13.13	0.592 ± 0.022	1504.71 \pm 5.73	21.84 \pm 2.28	0.0004 ± 10^{-4}	9340	14.14	0.365 ± 0.034	1389.11 \pm 307.17	53.25 \pm 0.70	0.015 ± 0.009
						2159.53 \pm 26.97	39.50 \pm 1.64					1917.25 \pm 698.18	51.06 \pm 1.37	
						4265.80 \pm 75.18	31.73 \pm 1.16					4090.16 \pm 506.84	30.88 \pm 2.71	
	100	bin(5)	2160	8.46	0.015 ± 0.009	1502.95 \pm 12.82	21.58 \pm 2.31	0.0004 ± 10^{-4}	2060	9.75	0.035 ± 0.013	1399.79 \pm 66.80	52.96 \pm 0.70	$0.0018 \pm 3 \times 10^{-4}$
						2158.97 \pm 31.45	39.35 \pm 1.64					2023.29 \pm 55.47	51.12 \pm 1.37	
						4266.46 \pm 62.45	31.56 \pm 1.18					4070.79 \pm 60.90	31.04 \pm 2.69	
0-100	10	unbin	4220	15.91	0.375 ± 0.034	1525.25 \pm 34.26	24.38 \pm 2.05	$2 \times 10^{-5} \pm 5 \times 10^{-5}$	9340	13.99	0.588 ± 0.022	1486.92 \pm 32.99	58.15 \pm 0.64	$7e-6 \pm 8 \times 10^{-5}$
						2156.31 \pm 150.12	43.25 \pm 1.49					2185.15 \pm 52.44	54.36 \pm 1.29	
						4256.61 \pm 323.22	34.87 \pm 1.06					4199.76 \pm 36.59	34.19 \pm 2.45	
	50	bin(5)	2180	9.75	0.020 ± 0.010	1528.14 \pm 50.61	24.55 \pm 2.03	$2 \times 10^{-5} \pm 5 \times 10^{-5}$	2070	9.88	0.028 ± 0.007	1392.78 \pm 46.47	58.15 \pm 0.64	$7e-6 \pm 8 \times 10^{-5}$
						2156.42 \pm 43.55	43.33 \pm 1.49					2025.80 \pm 28.99	56.29 \pm 1.25	
						4254.80 \pm 35.76	35.07 \pm 1.05					4072.25 \pm 57.81	34.65 \pm 2.41	
50	bin(5)	2180	9.75	0.020 ± 0.010	1516.75 \pm 9.35	10.88 \pm 4.59	$1 \times 10^{-4} \pm 5 \times 10^{-5}$	2080	7.32	0.025 ± 0.011	1497.67 \pm 2.82	32.51 \pm 1.15	$4e-4 \pm 8 \times 10^{-5}$	
					2157.42 \pm 379.77	19.34 \pm 3.35					2199.42 \pm 4.34	29.75 \pm 2.36		
					4272.82 \pm 143.82	15.59 \pm 2.36					4206.76 \pm 2.52	18.44 \pm 4.53		

The T_R and LRT statistics are computed for various binned and smoothed periodograms for larger time intervals. The frequencies in the table correspond with the centre of these fitted Lorentzian for the LRT test, and the maximum outlier of T_R , this is the QPO candidates. It also includes the fractional RMS amplitude (%) for the Lorentzians components fitted with the model (see Extended Data Fig. 5). Values that are referred to in the text are marked bold.

Article

Extended Data Table 3 | Z² search results and corresponding chance probabilities

Interval (Hz)	LED		HED	
	Peak Frequency (Hz)	Chance probability	Peak Frequency (Hz)	Chance probability
500 - 1100	835.9 ^{-84.7} _{+77.3}	1.2 x 10 ⁻⁴	-	-
1100 - 1700	1443.7 ^{-68.7} _{+74.8} ^a	4.9 x 10 ⁻²	1353.5 ^{-230.7} _{+217.7}	1.2 x 10 ⁻¹²
1800 - 2400	2131.7 ^{-151.0} _{+148.2}	2.4 x 10 ⁻⁹	2095.1 ^{-277.5} _{+180.8}	5.0 x 10 ⁻⁸
3900 - 4500	4249.7 ^{-102.7} _{+116.0}	1.7 x 10 ⁻⁴	4126.8 ^{-71.1} _{+73.0}	1.1 x 10 ⁻²

^aThe lower and upper bounds here mark the 95% confidence interval. The other bounds are the 99% confidence levels.